350-601

Number: 350-601
Passing Score: 800
Time Limit: 120 min
File Version: 1

350-601

Exam A

QUESTION 1
A user retrieves data in XML format from a Cisco APIC device by submitting a GET request on TCP port 443.
Which of the following technologies are most likely in use? (Choose two.)

A. JSON
B. REST API

C. HTTP
D. SOAP API
E. HTTPS

Correct Answer: BE
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Most likely, representational state transfer (REST) Application Programming Interface (API) and Hypertext Transfer Protocol Secure (HTTPS) are in use if a user retrieves data in Extensible Markup Language (XML) format from an application by submitting a GET request on Transmission Control Protocol (TCP) port 443. REST is an API architecture that uses Hypertext Transfer Protocol (HTTP) or HTTPS to enable external resources to access and make use of programming methods that are exposed by the API. In this way, users can interact with specific portions of a data structure from a remote system. By default, HTTPS operates on TCP port 443. A GET request is an HTTP method of retrieving information from an HTTP server.
It is not likely that HTTP is in use in this scenario, because the TCP port on which the GET request is being made is the HTTPS port. If unencrypted HTTP was being used in this scenario, the TCP port on which the request is being made would most likely be TCP port 80. By default, HTTP servers listen for traffic on TCP port 80. On Cisco Application Policy Infrastructure Controller (APIC) devices, HTTPS, not HTTP, is enabled by default. It is possible to enable HTTP on an APIC device. However, HTTP is less secure than HTTPS and is therefore not recommended for that purpose.

It is not likely that Simple Object Access Protocol (SOAP) API is being used in this scenario, because the user is retrieving data from a Cisco APIC device. Cisco APIC does not support SOAP, which is an older API messaging protocol that uses HTTP and XML to enable communication between devices. SOAP APIs are typically more resource-intensive than more modern APIs and, therefore, slower. Open APIs can be used to enable services such as billing automation and centralized management of cloud infrastructure.
JavaScript Object Notation (JSON) is not in use in this scenario. JSON is an output format that is supported by REST API. However, in this scenario, the user has retrieved data from the Cisco APIC device in XML format.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 15: Cloud Computing, Application Programming Interfaces, pp. 577-579
Cisco: HTTPS-HTTP Server and Client with SSL 3.0: Secure HTTP Server and Secure HTTP Client

QUESTION 2
During the APIC cluster discovery process, LLDP is used for which of the following tasks?

A. assignment of VTEP addresses
B. discovering MAC addresses
C. discovering private IP addresses
D. serving the APIC GUI

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Link-Layer Discovery Protocol (LLPD) is used for discovering private Internet Protocol (IP) addresses during the Cisco Application Policy Infrastructure Controller (APIC) cluster discovery process. LLDP is used by APIC controllers to discover the private IP addresses and other information assigned to other APIC controllers in the cluster. LLDP is a standard protocol that detects neighboring devices of any type.
LLDP is not used to assign virtual extensible local area network (VXLAN) tunnel endpoints (VTEPs). However, a Cisco Application Centric Infrastructure (ACI) fabric uses LLDP along with Dynamic Host Configuration Protocol (DHCP) to discover switch nodes and to assign IP addresses to VTEPs. LLDP is also used by APIC to detect virtual switches, although it is possible to use Cisco Discovery Protocol (CDP) for that purpose.
LLDP is not used to serve the APIC graphical user interface (GUI). The APIC GUI is constructed in Hypertext Markup Language (HTML) 5. Therefore, Hypertext Transfer Protocol (HTTP), not LLDP, serves the APIC GUI.
LLDP is not used to discover Media Access Control (MAC) addresses. A switched network uses Address Resolution Protocol (ARP) to map MAC addresses to IP addresses. Reference:

Cisco: Cisco Application Centric Infrastructure Fundamentals: Startup Discovery and Configuration

QUESTION 3
In which of the following formats does a REST API produce output? (Choose two.)

A. XML
B. CSV C. HTML
D. JSON
E. HTTP

Correct Answer: AD
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Representational state transfer (REST) produces output in either JavaScript Object Notation (JSON) or Extensible Markup Language (XML) format. REST is an Application Programming Interface (API) architecture that uses Hypertext Transfer Protocol (HTTP) or HTTP Secure (HTTPS) to enable external resources to access and make use of programming methods that are exposed by the API. For example, a web application that retrieves user product reviews from an online marketplace for display on third-party websites might obtain those reviews by using methods provided in an API that is developed and maintained by the marketplace. The JSON or XML output that is returned by the API is parsed by the third-party website for display.
A REST API does not produce output in Hypertext Markup Language (HTML) format. Although HTML is similar to XML, which uses tags like HTML does, XLM requires a strict syntax and is typically used to structure data, not format and render data in a web browser. HTML, on the other hand, is designed to inform a web browser about how given information should be displayed.
A REST API does not produce output in HTTP format. HTTP is the Open Systems Interconnection (OSI) networking model Application layer protocol that is used to transfer information from a web server to a web browser. A REST API uses HTTP to transmit requests for information to a web server, which is not the same as producing the formatted output that is returned from the server.
A REST API does not produce output in comma separated values (CSV) format. The CSV format is a common tabular format that is supported by spreadsheet applications and other business reporting applications. CSV files are plain-text files that segregate the fields of a table by using a combination of quotation marks, symbolic delimiters such as a comma or a semicolon, and line breaks.

Reference:

Cisco: Cisco APIC REST API Configuration Guide: About the REST API

QUESTION 4
You manage the Cisco ACI fabric in the following exhibit:

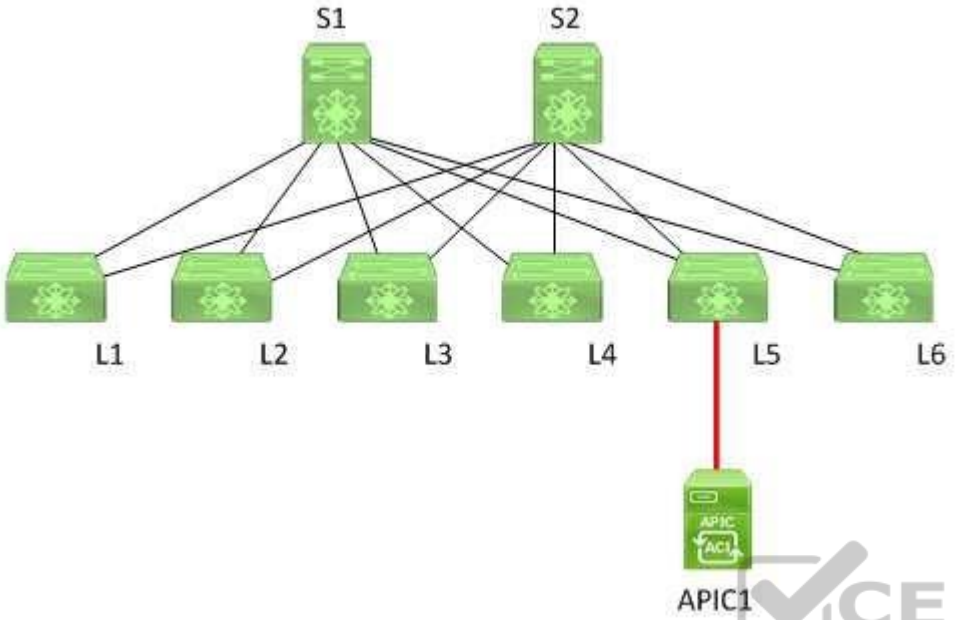

S2 has just been discovered by APIC1. Only one leaf switch has been discovered and registered with APIC1 so far. No other spine switches have been discovered. Which of the following switches will most likely be discovered next?

A. L5
B. S1
C. L1, L2, and L3
D. L4 and L6
E. L1, L2, L3, L4, and L6

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Most likely, S1 will be discovered next because of the way that Cisco Application Policy Infrastructure Controllers (APICs) discover and register leaf and spine switches in a Cisco Application Centric Infrastructure (ACI) fabric. In this scenario, APIC1 Has just discovered S2. This means that L5, the leaf switch to which APIC1 is directly connected, has already been discovered as well. L5 is directly connected to both spine switches in the topology. Because S2 was just discovered and no other spine switches have yet been discovered, it is likely that APIC1 will next discover S1.
APIC1 will not discover L5 next, because APIC1 has already discovered and registered L5. When a Cisco APIC begins the switch discovery process, it first detects only the leaf switch to which it is connected. After that leaf switch is registered, the APIC discovers each of the spine switches to which the leaf switch is connected. Spine switches do not automatically register with the APIC. When a spine switch is registered with the APIC, the APIC will discover all the leaf switches that are connected to that spine switch. Therefore, APIC1 will not discover all the leaf switches in this scenario until each spine switch is registered with the APIC. APIC1 will not discover L1, L2, L3, L4, or L6 next. All these switches are leaf switches that are directly connected to both S1 and S2. Although APIC1 in this scenario will eventually discover L1, L2, and L3, it will first discover all the spine switches to which L5 is directly connected.

Reference:

Cisco: Fabric Initialization and Switch Discovery: Switch Discovery

QUESTION 5
Which of the following is best described as the atomic units of work in Cisco UCS Director Orchestrator?

A. rollbacks
B. approvals
C. service requests
D. workflows
E. tasks

Correct Answer: E
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, tasks are best described as the atomic units of work in Cisco Unified Computing System (UCS) Director Orchestrator. Cisco UCS Director uses hardware abstraction to convent hardware and software into programmable actions that can then be combined into an automated custom workflow. Cisco UCS Director Orchestrator is the Cisco UCS Director engine that enables this automation. A task is a single action and is therefore the smallest unit of work. A workflow is a series of tasks and is therefore not the atomic unit of work in Cisco UCS Director Orchestrator. A workflow is a container that defines the order in which tasks should be performed. However, it is possible for a workflow to contain a single task. Workflows can be created and deployed from workflow templates. Service requests are created when a workflow is executed and are therefore not the atomic unit of work in Cisco UCS Director Orchestrator. Service requests are Cisco UCS Director processes that can be exist in one of several states. For example, a service request that has not run yet might exist in a scheduled state. A

service request that has been successfully executed exists in a completed state. A service request that was attempted but not successfully executed might exist in a failed state.
Approvals are points in a workflow that require user intervention. For example, a service request might exist in a blocked state if the request cannot complete until an administrator approves the service request. Approvals enable administrators to provide input values that can affect the product of a given workflow.
Rollbacks can be used to undo the results of workflows. For example, a workflow that creates unintended objects or components in a system can be rolled back so that those objects or components are removed. Cisco UCS Director Orchestrator rollbacks are different from relational database rollbacks in that they are not transactional. Instead, tasks in the workflow are executed in reserve order when a workflow is rolled back.

Reference:

Cisco: Orchestration and Automation: Cisco UCS Director Orchestrator

QUESTION 6
Which of the following are most likely to operate in the data plane of a Nexus switch? (Choose two.)

A. BGP
B. EIGRP
C. OSPF
D. store-and-forward switching
E. SNMP
F. cut-through switching

Correct Answer: DF
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, cut-through switching and store-and-forward switching are most likely to operate in the data plane of a Nexus switch. A Nexus switch consists of three operational planes: the data plane, which is also known as the forwarding plane, the control plane, and the management plane. Of the three, the data plane is where traffic forwarding occurs. Cut-through switching allows a switch to begin forwarding a frame before the frame has been received in its entirety. Store-and-forward switching receives an entire frame and stores it in memory before forwarding the frame to its destination.
Enhanced Interior Gateway Routing Protocol (EIGRP), Open Shortest Path First (OSPF), and Border Gateway Protocol (BGP) all operate in the control plane of a Nexus switch. The control plane is responsible for gathering and calculating the information required to make the decisions that the data plane needs for forwarding. Routing protocols operate in the control plane because they enable the collection and transfer of routing information between neighbors. This information is used to construct routing tables that he data plane can then use for forwarding.

Simple Network Management Protocol (SNMP) is an Internet Protocol (IP) network management protocol that operates in the management plane of a Nexus switch. The management plane is responsible for monitoring and configuration of the control plane. Therefore, network administrators typically interact directly with protocols running in the management plane.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 2: Management and Monitoring of Cisco Devices, Data Plane, pp. 50-53

QUESTION 7

You manage the Cisco FabricPath network in the following exhibit:

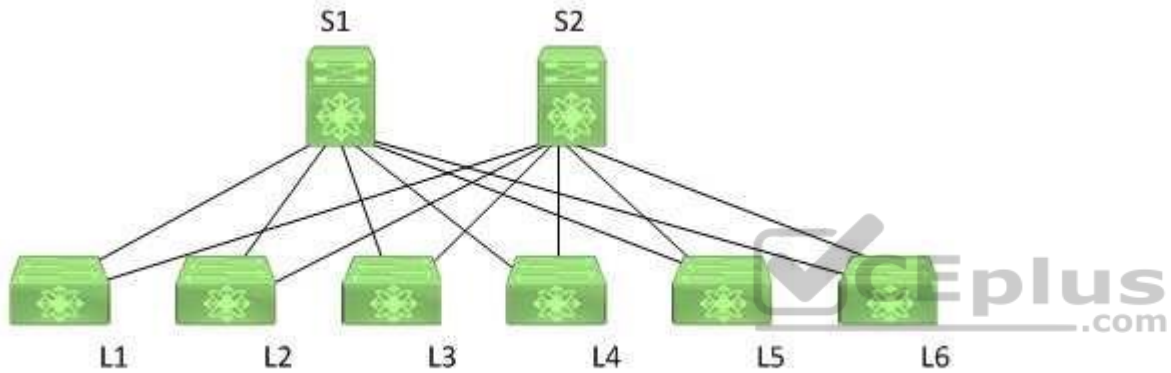

A host that is directly connected to L2 sends a unicast packet to a known host that is also directly connected to L2.
Which of the following is true?

A. L2 receives an unknown unicast packet from the sending host.
B. L2 will forward the packet to the multidestination tree.
C. L2 will not perform a lookup on the destination MAC address.
D. L2 will not forward the packet to a spine switch.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

L2 will not forward the packet to a spine switch. In this scenario, it is not necessary for L2 to forward the packet to a spine switch, because both the sending host and the receiving host are directly connected to L2. If the destination host in this scenario were connected to a different leaf switch, the known unicast packet would be forwarded from L2 to one of the spine switches and then on to the switch on which the destination host is located.
When L2 receives a known unicast packet, which is unicast packet with a destination address that the switch knows how to reach, that is intended for a host on a different switch the following occurs:

- L2 performs a Media Access Control (MAC) address lookup and located the destination switch.
- L2 looks up the switch ID of the destination switch and chooses a FabricPath core port on which to send the packet.
- L2 encapsulated the packet in a FabricPath header, including its own switch ID in the outer source address (OSA) field and the destination switch's ID in the outer destination address (ODA) field.
- L2 sends the packet.

By contrast, when L2 receives an unknown unicast packet, the following occurs:

- L2 updates the classic Ethernet MAC table with the sending host's source Mac address.
- L2 performs a lookup for the destination Mac address of the destination host.
- L2 encapsulates the classic Ethernet frame in a FabricPath header and floods it to the multidestination forwarding tree with L2's switch ID as the OSA. Because of the spine-and-leaf configuration of the topology, L2 would flood an unknown unicast packet to its directly connected spine switches, S1 and S2, which would then flood the packet to the other leaf switches. Each leaf switch would then perform a MAC address lookup. The switch that has the destination host in its MAC address table adds the sending host to its own MAC address table. At this point, the packet is sent to the destination host.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 1: Data Center Networking, FabricPath Packet Flow Example, pp. 38-40

QUESTION 8
Which of the following OTV technologies is a logical interface and cannot be a physical interface?

A. OTV join interface
B. OTV overlay interface
C. OTV internal interface
D. OTV edge device

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

An Overlay Transport Virtualization (OTV) overlay interface is a logical interface and cannot be a physical interface. OTV is a means of simplifying the deployment of data center interconnect (DCI), enabling the extension of Layer 2 applications between data centers. An OTV overlay interface is a logical interface that is

defined by the user. The OTV overlay interface receives and forwards any Open Systems Interconnection (OSI) networking model Layer 2 frames that must be transmitted to the remote site.
An OTV edge device is typically a virtual device context (VDC) running on A Cisco Nexus switch, such as a Nexus 7000 Series switch. The edge device is responsible for receiving Layer 2 traffic for virtual local area networks (VLANs) that are extended. Ethernet frame traffic is encapsulated into Internet Protocol (IP) packets are transmitted to the remote network.
An OTV join interface is an OTV technology that is a physical or logical Layer 3 interface. An OTV join interface is used to connect an overlay network to remote overlay edge devices. An OTV join interface can be a physical interface, subinterface, or logical interface such as a port channel.
An OTV internal interface is a Layer 2 interface. The OTV internal interface is an interface on the OTV edge device that connects to the VLANs that are being extended. Typically, the OTV internal interface operates similarly to any other Layer 2 switch trunk port or access port and does not require configuration specific to OTV.
Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 5: Data Center Overlay Networks, OTV Terminology, pp. 164-165

QUESTION 9
You are creating a workflow in UCS Director's Workflow Designer. You have connected each task's On Failure event in the workflow to the completed (failed) task. You need to connect each task's On Success event to the appropriate task.

Which of the following tasks are you most likely to choose?

A. start
B. completed (failed)
C. completed (success)
D. the next task in the workflow

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choice, you are most likely to connect each task's On Success event to the next task in the workflow. Workflows determine the order in which tasks that are designed to automate complex IT operations are performed. Workflow Designer allows administrators to create workflows that can be automated by using Unified Computing System (UCS) Director's orchestrator.
The following Cisco UCS Director' Workflow Designer tasks are predefined when creating a workflow: ▪
Completed (failed)
▪ Completed (success)
▪ Start

The start task is the beginning of the workflow. The completed (failed) task represents the end of a workflow when the desired result could not be achieved. The completed (success) task represent a successfully completed workflow. Each task in a workflow processes input and produces output that is sent to the next task in the workflow. In addition, each task contains an On Success event and an On Failure event that can be used to determine which task should be performed next based on whether the task could be successfully completed. On Success events should be connected to the next task in the workflow. On Failure events, on the other hand, should be connected to the completed (failed) task so that the workflow does not attempt to perform more tasks that would rely on successful output from the previously failed task.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 17: Understanding and Troubleshooting UCSD Workflows, Creating Workflows, pp. 641-645

QUESTION 10

Which of the following statements about Layer 3 virtualization on a Cisco Nexus 7000 Series switch is true?

A. A VRF represents a Layer 2 addressing domain.
B. Each VRF can belong to multiple VDCs.
C. Each Layer 3 interface is configured in one VRF.D. Each VDC can be configured with only one VRF.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Each Layer 3 interface on a Cisco Nexus 7000 Series switch can be configured with only one virtual routing and forwarding (VRF) instance. In other words, a Layer 3 interface that has been assigned to a given VRF cannot be simultaneously assigned to another VRF. There are two VRF instances configured on a Nexus 7000 Series switch by default: the management VRF and the default VRF. The management VRF is used only for management, includes only the Mgmt 0 interface, and uses only static routing. The default VRF, on the other hand, includes all Layer 3 interfaces until you assign those interfaces to another VRF.
A VRF instance represents an Open Systems Interconnection (OSI) networking model Layer 3 addressing domain, not a Layer 2 addressing domain. VRFs are used to logically separate OSI networking model Layer 3 networks. The address space, routing process, and forwarding table that are used within a VRF are local to that VRF.
Each VRF on a Cisco Nexus 7000 Series switch can belong to only one virtual device context (VDC). However, each VDC can be configured with multiple VRFs. A VDC is a virtual switch. Therefore, a VDC is a logical representation of a physical device on which VRFs can be configured.

Reference:

Cisco: Routing Overview: Layer 3 Virtualization

QUESTION 11

Which of the following examples best describes the SaaS service model?

A. A company obtains a subscription to use a service provider's infrastructure, programming tools, and programming languages to develop and serve cloud-based applications.
B. A company moves all company-wide policy documents to an internet-based virtual file system hosted by a service provider.
C. A company hires a service provider to deliver cloud-based processing and storage that will house multiple virtual hosts configured in a variety of ways.
D. A company licenses an office suite, including email service, that is delivered to the end user through a web browser.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A company that licenses an office suite, including email service, that is delivered to the end user through a web browser is an example of the Software as a Service (SaaS) service model. The National Institute of Standards and Technology (NIST) defines three service models in its definition of cloud computing: SaaS, Infrastructure as a Service (IaaS), and Platform as a Service (PaaS). Cloud computing offers several benefits over traditional physical infrastructure and software licensing, including a reduction in downtime and administrative overhead.
The SaaS service model enables its consumer to access applications running in the cloud infrastructure but does not enable the consumer to manage the cloud infrastructure or the configuration of the provided applications. A company that licenses a service provider's office suite and email service that is delivered to end users through a web browser is using SaaS. SaaS providers use an Internet-enabled licensing function, a streaming service, or a web application to provide end users with software that they might otherwise install and activate locally. Web-based email clients, such as Gmail and Outlook.com, are examples of SaaS. The PaaS service model provides its consumer with a bit more freedom than the SaaS model by enabling the consumer to install and possibly configure providersupported applications in the cloud infrastructure. A company that uses a service provider's infrastructure, programming tools, and programming languages to develop and serve cloud-based applications is using PaaS. PaaS enables a consumer to use the service provider's development tools or Application Programming Interface (API) to develop and deploy specific cloud-based applications or services. Another example of PaaS might be using a third party's MySQL database and Apache services to build a cloud-based customer relationship management (CRM) platform.
The IaaS service model provides the greatest degree of freedom by enabling its consumer to provision processing, memory, storage, and network resources within the cloud infrastructure. The IaaS service model also enables its consumer to install applications, including operating systems (OSs) and custom applications. However, with 1aaS, the cloud infrastructure remains in control of the service provider. A company that hires a service provider to deliver cloud-based processing and storage that will house multiple physical or virtual hosts configured in a variety of ways is using IaaS. For example, a company that wanted to establish a web server farm by configuring multiple Linux Apache MySQL PHP (LAMP) servers could save hardware costs by virtualizing the farm and using a provider's cloud service to deliver the physical infrastructure and bandwidth for the virtual farm. Control over the OS, software, and server configuration would remain the responsibility of the organization, whereas the physical infrastructure and bandwidth would be the responsibility of the service provider.
A company that moves all company-wide policy documents to an Internet-based virtual file system hosted by a third party is using cloud storage. Cloud storage is a term used to describe the use of a service provider's virtual file system as a document or file repository. Cloud storage enables an organization to conserve storage space on a local network. However, cloud storage is also a security risk in that the organization might not have ultimate control over who can access the files.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 15: Cloud Computing, Cloud Computing Service Models, pp. 574-577

QUESTION 12
You are installing a Cisco Nexus 1000v VSM in VMware vSphere by using the OVF folder method of installation.
Which of the following steps are you required to manually perform? (Choose four.)

A. perform initial VSM setup
B. add hosts
C. install VSM plug-in
D. configure VSM networking
E. create the VSM
F. configure SVS connection

Correct Answer: ABCF
Section: (none)
Explanation

Explanation/Reference:
Explanation:

You are required to manually perform the initial virtual supervisor module (VSM) setup, install the VSM plug-in, configure the server virtualization switch (SVS) connection, and add hosts. In this scenario, you are using the open virtualization format (OVF) folder method of installation. This method of installation automatically performs the first two steps of the six-step installation process. An OVF folder contains a hierarchy of files with different metadata that together define a given virtual environment.
There are three methods of installing the Nexus 1000v VSM: ▪
By using an International Standards Organization (ISO) image
▪ By using an OVF folder and installation wizard
▪ By using an open virtualization appliance (OVA) file and installation wizard
No matter which installation method is chosen, the VSM installation process consists of the following six steps: ▪
Creation of the VSM virtual machine (VM) in Cisco vCenter
▪ Configuration of VSM networking
▪ Initial VSM setup in the VSM console
▪ Installation of the VSM plug-in in vCenter
▪ Configuration of the SVS connection in the VSM console
▪ Addition of hosts to the virtual distributed switch in vCenter
The steps in this process that you are required to perform manually depend on the method of installation that you choose.
The OVA file method of installing the Nexus 1000v VSM is similar to the OVF folder method. However, the OVA file method performs the first four steps of the installation process. This means that you are required to manually configure the SVS connection in the VSM console and add hosts to the virtual distributed switch

in vCenter. However, the first four steps in the process are performed automatically after you deploy the OVA by using the installation wizard. The primary difference between an OVA file and an OVF folder is that the OVA file is a single compressed archive.
The ISO image method of installing the Nexus 1000v VSM requires that you manually perform each of the six steps in the installation process. An ISO image file contains a virtual filesystem that can be mounted by an operating system (OS) similar to an optical disc or a Universal Serial Bus (USB) flash drive.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 4: Cisco Nexus 1000V and Virtual Switching, Cisco Nexus 1000V VSM Installation Methods, pp. 144-145

QUESTION 13
Which of the following is true of the mgmt 0 interface?

A. Only this interface can be a member of the management VRF.
B. By default, it is assigned to the default VRF.
C. It can be assigned to any VRF.
D. It is used only by the default VDC.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Only the mgmt 0 interface can be a member of the management virtual routing and forwarding (VRF) instance. The management VRF is used only for management. No routing protocols are allowed to run in the management VRF, and all routing is static. The management VRF includes only the mgmt 0 interface, which cannot be assigned to any other VRF. However, the mgmt 0 interface is shared among virtual device contexts (VDCs).
VRFs are used to logically separate Open Systems Interconnection (OSI) networking model Layer 3 networks. Therefore, it is possible to have overlapping Internet Protocol version 4 (IPv4) or Internet Protocol version 6 (IPv6) addresses in environments that contain multiple tenants. However, an interface that has been assigned to a given VRF cannot be simultaneously assigned to another VRF. The address space, routing process, and forwarding table that are used within a VRF are local to that VRF. By default, a Cisco router is configured with two VRFs: the management VRF and the default VRF.
A Cisco router's default VRF instance is similar to a router's global routing table. The default VRF includes all Layer 3 interfaces until you assign those interfaces to another VRF. Similarly, the default VRF runs any routing protocols that are configured unless those routing protocols are assigned to another VRF. All **show** and **exec** commands that are issued in the default VRF apply to the default routing context.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 6: Virtualizing Cisco Network Devices, Describing Layer 3 Virtualization Within VDCs, pp. 204-206

QUESTION 14
Which of the following best describes the data plane of a Nexus switch?

A. It is where SNMP operates.
B. It is where routing calculations are made.
C. It is also known as the control plane.
D. It is where traffic forwarding occurs.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, the control plane of a Nexus switch is best described as where traffic forwarding occurs. A Nexus switch consists of three operational planes: the control plane, the management plane, and the data plane, which is also known as the forwarding plane. Cut-through switching allows a switch to begin forwarding a frame before the frame has been received in its entirety. Store-and-forward switching receives an entire frame and stores it in memory before forwarding the frame to its destination.
The control plane, not the data plane, of a Nexus switch is where routing calculations are made. Routing protocols such as Enhanced Interior Gateway Routing Protocol (EIGRP), Open Shortest Path First (OSPF), and Border Gateway Protocol (BGP) all operate in the control plane of a Nexus switch. The control plane is responsible for gathering and calculating the information required to make the decisions that the data plane needs for forwarding. Routing protocols operate in the control plane because they enable the collection and transfer of routing information between neighbors. This information is used to construct routing tables that the data plane can then use for forwarding.
The management plane, not the data plane, of a Nexus switch is where Simple Network Management Protocol (SNMP) operates. The management plane is responsible for monitoring and configuration of the control plane. Therefore, network administrators typically interact directly with protocols running in the management plane.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 2: Management and Monitoring of Cisco Nexus Devices, Operational Planes of a

Nexus Switch, pp. 50-59

QUESTION 15
You are configuring port mirroring on a Cisco switch. You configure a VLAN as the source port. You configure a physical Ethernet port as the destination port. Which of the following are you most likely configuring? (Choose two.)

A. VSPAN

B. SPAN
C. RSPAN
D. ERSPAN
E. RAP

Correct Answer: AB
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Most likely, you are configuring Switched Port Analyzer (SPAN) and virtual local area network (VLAN)-based SPAN (VSPAN) if you enable port mirroring by configuring a VLAN as the source port and a physical Ethernet port as the destination port on the same Cisco switch. SPAN source and destination ports must be on the same device. SPAN is limited to monitoring traffic on only the local device and cannot direct traffic to destination ports on a separate device for analysis. The source port can be a physical or virtual Ethernet port, a port channel, or a VLAN if VSPAN is being used. The destination port can be a physical or virtual Ethernet port or a port channel. However, the source port and the destination port cannot be the same port. All ports in a source VLAN become SPAN source ports. SPAN, Remote SPAN (RSPAN), and Encapsulated RSPAN (ERSPAN) are all capable of using VLANs as sources by implementing VSPAN.
You are not configuring RSPAN in this scenario. RSPAN source and destination ports must be on the same local area network (LAN) but are not restricted to the same device. RSPAN enables you to monitor traffic on a network by capturing and sending traffic from a source port on one device to a destination port on a different device on a nonrouted network.
You are not configuring ERSPAN in this scenario. ERSPAN source and destination ports can exist in different LANs and are not typically configured on the same switch. ERSPAN enables an administrator to capture and analyze traffic across a routed network. Therefore, ERSPAN can monitor traffic across multiple routers on a network that spans multiple locations.
You are not configuring Roving Analysis Port (RAP) in this scenario. RAP is a form of port mirroring that is configured on 3Com switches, not on Cisco devices.

Reference:

Cisco: Configuring Local SPAN, Remote SPAN (RSPAN), and Encapsulated RSPAN (ERSPAN): RSPAN Overview

QUESTION 16
Examine the Cisco UCS Director Workflow Designer in the following exhibit:

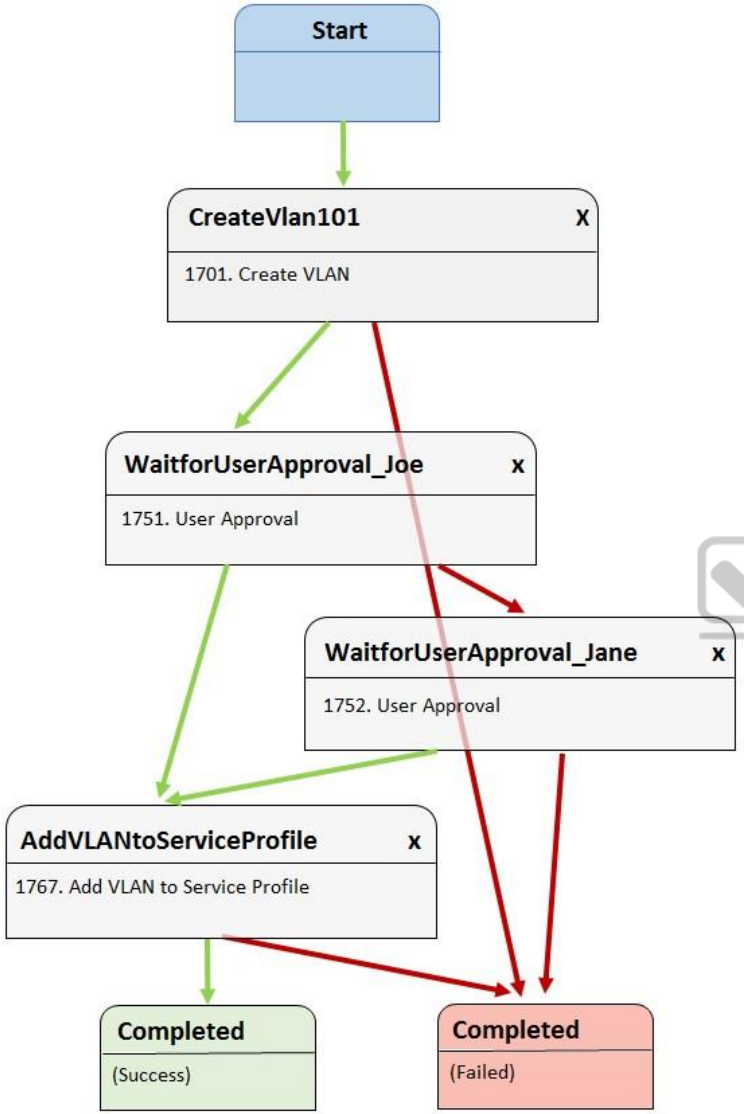

Joe initiates the workflow and approves the addition of the VLAN to the service profile.
Which of the following best describes what will occur next?

A. VLAN 101 will be added to the service profile if Jane approves.
B. VLAN 101 will be successfully added to the service profile.
C. VLAN 101 will not be added to the service profile, and the workflow will succeed.
D. VLAN 101 will not be added to the service profile, and the workflow will fail.
E. VLAN 101 will be sent back to Joe for approval if Jane approves.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Most likely, virtual local area network (VLAN) 101 will be added to the service profile if Joe approves. The Cisco Unified Computing System (UCS) Director Workflow Designer workflow in this scenario creates VLAN 101 after it starts. Before the next task is executed, the workflow requires the approval of a user named Joe. In this scenario, Joe has approved this step. Because Joe's approval is obtained, the workflow continues to the **AddVLANtoServiceProfile** task.
Cisco UCS Director Workflow Designer is a graphical user interface (GUI) that enables users to create automated workflows in a drag-and-drop fashion. Each task in a workflow is equipped with an **On Success** button and an **On Failure** button. Each button provides a drop-down list of other tasks in the workflow. In this way, the user can select which tasks are executed next if a task succeeds and which tasks are executed next if a task fails. Green arrows in Workflow Designer represent the **On Success** path. Red arrows represent the **On Failure** path.
Jane will not be provided with an opportunity to approve or reject in this scenario, because Joe has already approved the workflow. The **On Success** path of the **WaitforUserApproval_Joe** task is directly tied to the **AddVLANtoServiceProfile** task. If Joe had not approved, Jane would have been provided the opportunity to approve or reject because the **On Failure** path of the **WaitforUserApproval_Joe** task is tied to the **WaitforUserApproval_Jane** task. In that case, Jane could approve, which would launch the **AddVLANtoServiceProfile** task, or reject, which would cause the workflow to end in failure.
There is nothing in this scenario to indicate whether the **AddVLANtoServiceProfile** task would be successful following Joe's approval. Therefore, it is not possible to determine whether the workflow will follow that task's On Success path or its **On Failure** path.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 17: Understanding and Troubleshooting UCSD Workflows, Creating Workflows, pp. 641-645

QUESTION 17
Which of the following is not a physical device?

A. Nexus 5000 Series
B. Nexus 100v Series
C. Nexus 9000 Series
D. Nexus 2000 Series
E. Nexus 7000 Series

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, the Cisco Nexus 1000v Series is not a physical device. The Cisco Nexus 1000v is a virtual switch that is capable of connecting to upstream physical switches in order to provide connectivity for a virtual machine (VM) network environment. Although the Cisco Nexus 1000v operates similar to a standard switch, it exists only as software in a virtual environment and is therefore not a physical switch.
The Cisco Nexus 2000 Series of switches are physical devices. However, they are fabric extenders (FEXs) and cannot operate as standalone switches. FEX technologies depend on parent switches, such as a Cisco Nexus 5500 Series switch or a Cisco Nexus 7000 Series switch, to provide forwarding tables and control plane functionality. FEX technologies are intended to extend the network to edge devices. Typically, FEX devices in the Cisco Nexus 2000 Series are managed by first connecting to the parent device by using either Telnet or Secure Shell (SSH) and then configuring the FEX.
Cisco Nexus 5000 Series switches are physical devices that operate as standalone switches. Cisco Nexus 5000 Series switches are data center access layer switches that can support 10-gigabit-per-second (Gbps) or 40-Gbps Ethernet, depending on the model. Native Fibre Channel (FC) and FC over Ethernet (FCoE) are also supported by Cisco Nexus 5000 Series switches.
Cisco Nexus 7000 Series switches are physical devices that operate as standalone switches. Cisco Nexus 7000 Series switches are typically used as an end-toend data center solution, which means that the series is capable of supporting all three layers of the data center architecture: core layer, aggregation layer, and access layer. In addition, the Cisco Nexus 7000 Series supports virtual device contexts (VDCs). The Cisco Nexus 7000 Series can support up to 100-Gbps Ethernet.
Cisco Nexus 9000 Series switches are physical devices that operate as standalone switches. Cisco Nexus 9000 Series switches can operate either as traditional NX-OS switches or in an Application Centric Infrastructure (ACI) mode. Unlike Cisco Nexus 7000 Series, Cisco Nexus 9000 Series switches do not support VDCs or storage protocols.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 4: Cisco Nexus 1000V and Virtual Switching, Cisco Nexus 1000V System Overview, p. 134

QUESTION 18
Which of the following Cisco Unified Fabric Features can result in reduced cabling?

A. VXLANs
B. vPCs
C. STP
D. CNAs

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, converged network adapters (CNAs) are the Cisco Unified Fabric feature that can result in reduced cabling. Cisco Unified Fabric is a combination of architecture and high performance concepts that is intended to simplify data center networks. CNAs are network adapters that combine network interface cards (NICs) and host bus adapters (HBAs), enabling one adapter to support both Ethernet and Fibre Channel (FC). A server that contains separate FC and local area network (LAN) ports can require significantly more cabling than a server that is configured with a CNA. Cisco Unified Fabric also helps converge a data center's LAN and storage area network (SAN) over a single transport in order to simplify management, provisioning, and operation.
Virtual Port Channels (vPCs) do not necessarily result in reduced cabling. A vPC is a Cisco Unified Fabric alternative to a traditional EtherChannel port channel. Therefore, vPCs are intended to create high-bandwidth redundant links between Layer 2 devices. Traditional EtherChannel relies on Spanning Tree Protocol (STP).
Port channels that are created by using vPCs still rely on STP to mitigate switching links if they occur, but do not rely on it in the same functional way that EtherChannel does.
Virtual extensible LANs (VXLANs) are not a Cisco Unified Fabric that can result in reduced cabling. VXLANs use Layer 3 technologies to extend Layer 2 technologies. In this way, VXLANs can achieve the same results as Cisco FabricPath without implementing FabricPath. Cisco FabricPath enables the scaling of a Layer 2 network beyond normal practical limitations by using Layer 3 routing protocol Intermediate System-to-Intermediate System (IS-IS) in place of STP.
STP is not a Cisco Unified Fabric feature. Instead, STP is a Layer 2 protocol that is intended to prevent switching loops in a network with redundant links.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 3: Unified Fabric Overview, Convergence of Network and Storage, pp. 82-88

QUESTION 19
Which of the following frame fields does Cisco FabricPath use to mitigate temporary Layer 2 loops?

A. FTAG
B. LID
C. STP
D. TTL
E. IS-IS

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, the Time to Live (TTL) field is the Cisco FabricPath frame field that FabricPath uses to mitigate temporary Open Systems Interconnection (OSI) networking model Layer 2 loops. The TTL field in a FabricPath frame operates similarly to the TTL field in Internet Protocol (IP) networking in that the field is decremented by a value of 1 each time it traverses a new hop. If the TTL expires, the frame is discarded.
Cisco FabricPath frames are classic Ethernet frames that are encapsulated with a 16-byte FabricPath header. This header contains a 48-bit outer destination address (ODA), which is a Media Access Control (MAC) address, and 48-bit outer source address (OSA). In addition, the field contains a 32-bit FabricPath tag. The classic Ethernet frame's cyclic redundancy check (CRC) field is replaced by a new CRC field that is updated to reflect the additional header data in the frame. The TTL field is a 6-bit field that resides at the end of the FabricPath tag field.
Cisco FabricPath does not use an Intermediate System-to-Intermediate System (IS-IS) frame field to mitigate temporary Layer 2 loops. In addition, FabricPath does not use a Spanning Tree Protocol (STP) frame field. However, the IS-IS routing protocol is used as a Layer 3 replacement for traditional STP in a Cisco FabricPath topology. In traditional networking, STP is used to prevent Layer 2 switching loops in a topology that contains redundant links. The use of the IS-IS routing protocol ensures that Cisco FabricPath operates as a multipath environment for Layer 2 packets. In other words, IS-IS ensures that Cisco FabricPath is capable of Layer 2 multipath forwarding.
Cisco FabricPath does not use the forwarding tag (FTAG) frame field to mitigate temporary Layer 2 loops. Instead, the FTAG field is used to identify the unique FabricPath topology that unicast traffic is traversing. Each topology in FabricPath is assigned a unique tag. For multicast or broadcast traffic, the 10-bit FTAG field contains an ID for a forwarding tree that contains multiple destinations within the topology. The FTAG field is the second of three fields that reside in the FabricPath tag field. It is preceded by the 16-bit Ethertype field and succeeded by the TTL field.
Cisco FabricPath does not use the local ID (LID) frame field to mitigate temporary Layer 2 loops. Instead, the LID field stores a 16-bit value that identifies the edge port that a packet is either destined to or sent from. The LID field is the last field in the ODA and OSA fields of a Cisco FabricPath header. The edge port can be either a physical port or a logical port. In addition, the value in the LID field is locally significant to the switch that the frame is traversing.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 1: Data Center Networking, FabricPath Frame Format, pp. 34-35

QUESTION 20
Which of the following are ways that Cisco VIC adapters can communicate? (Choose three.)

A. by using cut-through switching
B. by using store-and-forward switching
C. by using a bare-metal OS driver
D. by using pass-through switching
E. by using software and a Nexus 1000v

Correct Answer: CDE
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Cisco virtual interface card (VIC) adapters support three communication methods: by using software and a Cisco Nexus 1000v virtual switch, by using pass-through switching, and by using a bare-metal operating system (OS) driver. Cisco VICs integrate with virtualized environments to enable the creation of virtual network interface cards (vNICs) in virtual machines (VMs). Traffic between VMs is controlled by the hypervisor when using software and a Cisco Nexus 1000v virtual switch.
Cisco describes the software-based method of handling traffic by using a Cisco Nexus 1000v virtual switch as Virtual Network Link (VN-Link).
Pass-through switching, which Cisco describes as hardware-based VN-Link, is a faster and more efficient means for Cisco VIC adapters to handle traffic between VMs. Pass-through switching uses application-specific integrated circuit (ASIC) hardware switching, which reduces overhead because the switching occurs in the fabric instead of relying on software. Pass-through switching also enables administrators to apply network policies between VMs in a fashion similar to how traffic policies are applied between traditional physical network devices.
Cisco VIC adapters forward traffic similar to other Cisco Unified Computing System (UCS) adapters when installed in a server that is configured with a single baremetal OS without virtualization. It is therefore possible to use a Cisco VIC adapter with OS drivers to create static vNICs on a server in a nonvirtualized environment. Cisco VIC adapters do not communicate by using store-and-forward switching. However, switches can be configured to use store-and-forward switching. A switch that uses store-and-forward switching receives the entire frame before forwarding the frame. By receiving the entire frame, the switch can verify that no cyclic redundancy check (CRC) errors are present in the frame; this helps prevent the forwarding of frames with errors. Cisco VIC adapters are hardware interface components that support virtualized network environments.
Cisco VIC adapters do not communicate by using cut-through switching. The cut-through switching method is a switch forwarding method that begins forwarding a frame as soon as the frame's destination address is received, which is after the first six bytes of the packet are copied. Thus the switch begins forwarding the frame before the frame is fully received, which helps reduce latency. However, with the cut-through method, the frame is not checked for errors prior to being forwarded.

Reference:

Cisco: Overview of VN-Link in Cisco UCS

QUESTION 21
You are configuring a service profile for a Cisco UCS server that has been configured with two converged network adapters.
What is the maximum number of vHBAs that can be configured on this server if no vNICs are configured?

A. four
B. six
C. eight
D. two

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

If no virtual network interface cards (vNICs) are configured on the Cisco Unified Computing System (UCS) server in this scenario, up to four virtual host bus adapters (vHBAs) could be configured. A converged network adapter is a single unit that combines a physical host bus adapter (HBA) and a physical Ethernet network interface card (NIC). A Cisco converged network adapter typically contains two of these types of ports. Therefore, if no vNICs are configured on the device, it is possible for the Cisco UCS server in this scenario to be configured with four vHBAs.
Although it is possible to configure two vHBAs on the Cisco UCS server in this scenario, it is possible to configure a maximum of four vHBAs because no vNICs have been configured. If two vNICs had been configured in this scenario, it would be possible to configure a maximum of two vHBAs.
It is not possible to configure six or eight vHBAs, because the Cisco UCS server in this scenario is not configured with more than two converged network adapters. In order to support a higher maximum number of vHBAs, the Cisco UCS server in this scenario would need to be configured with a greater number of physical HBAs or converged network adapters.

Reference:

Cisco: Overview of Cisco Unified Computing System: Configuration through Service Profiles

QUESTION 22
Which of the following VMware ESXi VM migration solutions are capable of migrating VMs without first powering off or suspending them? (Choose two.)

A. cloning
B. cold migration
C. Storage vMotion
D. copying

E. vSphere vMotion

Correct Answer: CE
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, only VMware vSphere vMotion and Storage vMotion are capable of migrating VMware ESXi virtual machines (VMs) without first powering off or suspending them. VMware's ESXi server is a bare-metal server virtualization technology, which means that ESXi is installed directly on the hardware it is virtualizing instead of running on top of another operating system (OS). This layer of hardware abstraction enables tools like vMotion to migrate ESXi VMs from one host to another without powering off the VM, enabling the VM's users to continue working without interruption.
VMware's vSphere vMotion can be used to migrate a VM's virtual environment from one host to another. However, vSphere vMotion does not migrate datastores. Storage vMotion, on the other hand, allows migration of both the VM and its datastore. The datastore is the repository of VM-related files, such as logs and virtual disks. When migrating a VM by using Storage vMotion, both the virtualized environment and the datastore can be moved to a new host without powering down or suspending the VM.
Cold migration cannot migrate a VM without powering down or suspending the VM. Cold migration is the process of powering down a VM and moving the VM or the VM and its datastore to a new location. While a cold migration is in progress, no users can perform tasks inside the VM.
Neither copying nor cloning can migrate a VM. Both copying and cloning create new instances of a given VM. Therefore, neither action is a form of migrating a VM to another host. Typically, a VM must be powered off or suspended in order to successfully copy or clone it.

Reference:

VMware: VMware Docs: Migrating Virtual Machines
VMware: Virtualization Overview (PDF)

QUESTION 23
You are configuring a Cisco Nexus 5000 Series switch for the first time. The switch requires an IPv4 management interface.
Which of the following setup information is optional at first boot? (Choose three.)

A. additional account names and passwords for account creation
B. an SNMP community string
C. a switch name
D. an IPv4 subnet mask for the management interface
E. a new strong password for the admin user

Correct Answer: ABC

Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, a switch name, additional account names and passwords for account creation, and a Simple Network Management Protocol (SNMP) community string are all optional when configuring a Cisco Nexus 5000 Series switch for the first time if the switch requires an Internet Protocol version 4 (IPv4) management interface. If you configure a switch name, the name you choose for the switch will also be used at the command-line interface (CLI) prompt. If you choose not to configure a switch name during initial setup, the default name for the switch is **switch**.
Configuring additional account names and passwords when you are configuring a Cisco Nexus 5000 Series switch for the first time is optional. At first boot, a Cisco Nexus 5000 Series switch is configured with the single user account named admin. This account is also the network admin and cannot be changed or deleted.
Configuring an SNMP community string when you are configuring a Cisco Nexus 5000 Series switch for the first time is optional. An SNMP read-only community name enables another SNMP device with that community string to request SNMP management information from the Nexus switch.
Only a new strong password for the admin user and an IPv4 subnet mask for the management interface are required. In this scenario, you are configuring a Cisco Nexus 5000 Series switch for the first time. The switch requires an IPv4 management interface. Because the switch requires an IPv4 management interface, a subnet mask for the management interface must be configured at first boot.
When configuring a Cisco Nexus 5000 Series switch for the first time, you will be required to configure an admin password before configuration. This step in the configuration process is required and cannot be skipped by using the Ctrl-C keyboard combination. After you have successfully configured an admin password, you can enter setup mode by entering **yes** at the prompt.

Reference:

Cisco: Initial Switch Configuration: Preparing to Configure the Switch

QUESTION 24
Which of the following is not an advantage of using Cisco UCS identity pools with service profile templates?

A. Many templates can be updated.
B. Identities are manually assigned.
C. Speed and flexibility in server creation are increased.
D. Many service profiles can be updated.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Identities do not need to be manually assigned when using Cisco Unified Computing System (UCS) identity pools with service profile templates. Identity pools are logical resource pools that can be read and consumed by a service profile or a service profile template. These pools are used to uniquely identify groups of servers that share the same characteristics. A service profile can be used from Cisco UCS Manager to automatically apply configurations to the servers identified by the pool.
For Cisco UCS configurations or scenarios that require virtualized identities, the use of identity pools can greatly speed the server creation and template updating processes. There are six types of identity pools that can serve Cisco UCS service profile templates: Internet Protocol (IP) pools
▪ Media Access Control (MAC) pools
▪ Universally unique identifier (UUID) pools ▪
World Wide Node Name (WWNN) pools
▪ World Wide Port Name (WWPN) pools
▪ World Wide Node/Port Name (WWxN) pools
IP identity pools contain IP addresses, which are 32-bit decimal addresses that are assigned to interfaces. In a Cisco UCS domain, IP pools are typically used to assign one or more management IP addresses to each server's Cisco Integrated Management Controller (IMC).
MAC identity pools contain MAC addresses, which are 48-bit hexadecimal addresses that are typically burned into a network interface card (NIC). The first 24 bits of a MAC address represent the Organizationally Unique Identifier (OUI), which is a value that is assigned by the Institute of Electrical and Electronics Engineers (IEEE). The OUI identifies the NIC's manufacturer. The last 24 bits of a MAC address uniquely identify a specific NIC constructed by the manufacturer. This value is almost always an identifier that the manufacturer has never before used in combination with the OUI.
UUID identity pools contain Open Software Foundation 128-bit addresses. These addresses, known as UUIDs, contain a prefix and a suffix. The prefix identifies the unique UCS domain. The suffix is assigned sequentially and can represent the domain ID and host ID. UUIDs are typically used to assign software licenses to a given device.
The WWNN identity pool is a single pool for an entire Cisco UCS domain. WWNNs are 64-bit globally unique identifiers that specify a given Fibre Channel (FC) node. These identifiers are typically used in storage area network (SAN) routing.
Similar to the WWNN identity pool, the WWPN identity pool contains globally unique 64-bit identifiers. However, WWPNs represent a specific FC port, not an entire node.
WWSN identity pools contain a mix of WWNNs and WWPNs. WWxN pools can be used in any place in Cisco UCS Manager that can use WWNN pools and WWPN pools.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 9: Cisco Unified Computing System Pools, Policies, Templates, and Service Profiles, Cisco UCS Logical Resource Pools, pp. 340-341

QUESTION 25
You connect Port 12 of a Cisco UCS Fabric Interconnect to a LAN. All the Fabric Interconnect ports are unified ports.
Which of the following options should you select for Port 12 in Cisco UCS Manager?

A. **Configure as Server Port**
B. **Configure as FCoE Storage Port**
C. **Configure as FCoE Uplink Port**

D. **Configure as Uplink Port**

E. **Configure as Appliance Port**

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

You should select the Cisco Unified Computing System (UCS) Manager **Configure as Uplink Port** option for Port 12 on the Cisco UCS Fabric Interconnect in this scenario because Port 12 is connected to a local area network (LAN). Cisco UCS Fabric Interconnects enable management of two different types of network fabrics in a single UCS domain, such as a Fibre Channel (FC) storage area network (SAN) and an Ethernet LAN. The **Configure as Uplink Port** option configures Port 12 in the uplink role in Ethernet mode. This port role is used for connecting a UCS Fabric Interconnect to an Ethernet LAN.
Unified ports in a UCS Fabric Interconnect can operate in one of two primary modes: FC or Ethernet. Each mode can be assigned a different port role, such as server port, uplink port, or appliance port for Ethernet. The FC port mode supports two roles: FC over Ethernet (FCoE) uplink port and FCoE storage port. Uplink ports are used to connect the Fabric Interconnect to the next layer of the network, such as to a SAN by using the FCoE uplink port role or to a LAN by using the uplink port role.
You should not choose the **Configure as FCoE Uplink Port** option in this scenario. The **Configure as FCoE Uplink Port** option is used to connect the Fabric Interconnect to an FC SAN.
You should not choose the **Configure as Server Port** option in this scenario. The **Configure as Server Port** option is used to connect the Fabric Interconnect to network adapters on server hosts.
You should not choose the **Configure as FCoE Storage Port** option in this scenario. The **Configure as FCoE Storage Port** option is used to connect the Fabric Interconnect to a Direct-Attached Storage (DAS) device.
You should not choose the **Configure as Appliance Port** option in this scenario. The **Configure as Appliance Port** option is used to connect the Fabric Interconnect to a storage appliance.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 8: Cisco UCS Manager, Basic Port Roles in the Cisco UCS Fabric Interconnects, pp. 318-319

Cisco: Configuring Ports: Server and Uplink Ports on the Fabric Interconnect

QUESTION 26
You are configuring a vPC domain on SwitchA, SwitchB, and SwitchC. You connect all three switches by using 10-Gbps Ethernet ports. Next, you issue the following commands on each switch, in order:
feature vpc vpc
domain 101

You now want to configure the peer keepalive links between the switches.
Which of the following is true?

A. You will not be able to peer any of the switches.
B. You can peer all three switches if you add a vPC domain ID to SwitchB.
C. You will be able to peer only two of the three switches.
D. You can peer all three switches if you change the vPC domain ID.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

You will be able to peer only two of the three switches in this scenario because it is not possible to configure more than two switches per virtual port channel (vPC) domain. A vPC domain is comprised of two switches per domain. Each switch in the vPC domain must be configured with the same vPC domain ID. To enable vPC configuration on a Cisco Nexus 7000 Series switch, you should issue the **feature vpc** command on both switches. To assign the vPC domain ID, you should issue the **vpc domain** *domain-id* command, where *domain-id* is an integer in the range from 1 through 1000, in global configuration mode. For example, issuing the **vpc domain 101** command on a Cisco Nexus 7000 Series switch configures the switch with a vPC domain ID of 101.
A vPC forms an Open Systems Interconnection (OST) networking model Layer 2 port channel. Each virtual device context (VDC) is a separate switch in a vPC domain. A VDC logically virtualizes a switch. A VDC is a single virtual instance of physical switch hardware. A vPC logically combines ports from multiple switches into a single port-channel bundle. Conventional port channels, which are typically used to create high-bandwidth trunk links between two switches, require that all members of the bundle exist on the same switch. vPCs enable virtual domains that are comprised of multiple physical switches to connect as a single entity to a fabric extender, server, or other device.
Only one vPC domain can be configured per switch. If you were to issue more than one **vpc domain** *domain-id* command on a Cisco Nexus 7000 Series switch, the vPC domain ID of the switch would become whatever value was issued last. After you issue the **vpc domain** *domain-id* command, the switch is placed into vPC domain configuration mode. In vPC domain configuration mode, you should configure a peer keepalive link.
Peer keepalive links monitor the remote device to ensure that it is operational. You can configure a peer keepalive link in any virtual routing and forwarding (VRF) instance on the switch. Each switch must use its own Internet Protocol (IP) address as the peer keepalive link source IP address and the remote switch's IP address as the peer keepalive link destination IP address. The following commands configure a peer keepalive link between SwitchA and SwitchB in vPC domain 101:

```
SwitchA(config)#vpc domain 101
SwitchA(config-vpc-domain)#peer-keepalive destination 192.168.1.2 source 192.168.1.1 vrf default
SwitchB(config)#vpc domain 101
SwitchB(config-vpc-domain)#peer-keepalive destination 192.168.1.1 source 192.168.1.2 vrf default
```

A vPC peer link should always be a 10-gigabit-per-second (Gbps) Ethernet port. Peer links are configured as a port channel between the two members of the vPC domain. You should configure vPC peer links after you have successfully configured a peer keepalive link. Cisco recommends connecting two 10-Gbps Ethernet

ports from two different input/output (I/O) modules. To configure a peer link, you should issue the **vpc peer-link** command in interface configuration mode. For example, the following commands configure a peer link on Port-channel 1:

```
SwitchA(config)#interface port-channel 1
SwitchA(config-if)#switchport mode trunk SwitchA(config-
if)#vpc peer-link
SwitchB(config)#interface port-channel 1
SwitchB(config-if)#switchport mode trunk
SwitchB(config-if)#vpc peer-link
```

It is important to issue the correct **channel-group** commands on a port channel's member ports prior to configuring the port channel. For example, if you are creating Port-channel 1 by using the Ethernet 2/1 and Ethernet 2/2 interfaces, you could issue the following commands on each switch to correctly configure those interfaces as members of the port channel:

```
SwitchA(config)#interface range ethernet 2/1-2
Switch(config-if-range)#switchport
SwitchA(config-if-range)#channel-group 1 mode active
SwitchB(config)#interface range ethernet 2/1-2
SwitchB(config-if-range)#switchport
SwitchB(config-if-range)#channel-group 1 mode active
```

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 1: Data Center Networking, vPC Limitations, p. 27

QUESTION 27
Which of the following Cisco UCS servers are comprised of cartridge-style modules?

A. B-Series servers
B. UCS Mini servers
C. E-Series servers
D. C-Series servers
E. M-Series servers

Correct Answer: E
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, only Cisco Unified Computing System (UCS) M-Series servers are comprised of cartridge-style modules. Cisco UCS M-Series servers are neither rack servers nor blade servers. M-Series servers are comprised of computing modules that are inserted into a modular chassis. The chassis, not the server, can be installed in a rack. The modular design enables the segregation of computing components from infrastructure components.
Cisco UCS B-Series servers, Cisco UCS E-Series servers, and Cisco UCS Mini servers are all blade servers that typically reside in a blade chassis. UCS B-Series blade servers can connect only to Cisco UCS Fabric Interconnect, not directly to a traditional Ethernet network. Blade servers in a chassis are typically hotswappable, unlike the components of a rack-mount server. Therefore, blade server configurations are less likely to result in prolonged downtime if hardware fails. Cisco UCS E-Series servers are typically installed in a blade chassis. However, Cisco UCS E-Series servers have similar capabilities to the standalone C-Series servers and do not require connectivity to a UCS fabric. In a small office environment, Cisco UCS E-Series servers are capable of providing the network connectivity and capabilities of a C-Series server along with the availability of B-Series servers.
Cisco UCS Mini servers are a compact integration of a Cisco UCS 5108 blade chassis, Cisco UCS 6324 Fabric Interconnects, and Cisco UCS Manager. Unlike other Cisco UCS servers, the UCS Mini server requires a special version of Cisco UCS Manager for management. It is possible to connect Cisco UCS C-Series servers to Cisco UCS Mini servers in order to expand their abilities.
Cisco UCS C-Series servers are typically installed directly in a rack. Cisco UCS C-Series servers are rack-mount standalone servers that can operate either with or without integration with Cisco UCS Manager. Therefore, Cisco UCS C-Series servers do not require a UCS fabric. For administrators who are more familiar with traditional Ethernet networks than UCS Fabric Interconnect, C-Series servers will most likely be simpler to deploy and feel more familiar than other Cisco UCS server products.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 7: Cisco UCS Architecture, Cisco UCS M-Series Servers, pp. 272-273

QUESTION 28
Which of the following devices do not directly connect to a spine switch? (Choose four.)

A. another spine switch
B. a router
C. a leaf switch
D. a server
E. an APIC controller

Correct Answer: ABDE
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, only a leaf switch can connect to a spine switch. In a spine-leaf architecture, the leaf layer of switches provides connectivity to and scalability for all other devices in the data center network. However, leaf switches do not connect to other leaf switches. Instead, a leaf switch communicates with

another leaf switch by using a spine switch. The following exhibit displays a typical spine-leaf architecture wherein the top row of devices represents the spine layer of switches and the bottom row of devices represents the leaf layer of switches:

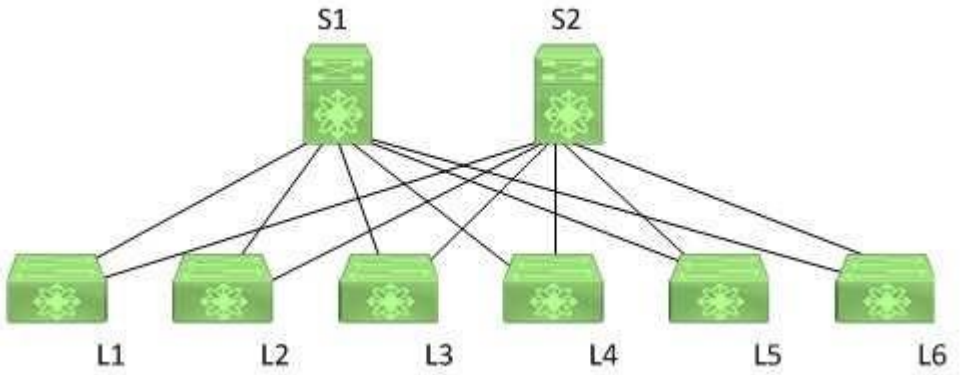

In the exhibit above, the leaf switches are each directly connected to both spine switches. Switches S1 and S2 comprise the spine layer of the topology. Switches L1, L2, L3, L4, L5, and L6 comprise the leaf layer of the topology. In order for L1 to send a packet to L6, the packet must traverse either S1 or S2. The spine-leaf architecture differs from the traditional three-tier network architecture, which consists of a core layer, an aggregation layer, and an access layer.
Leaf switches connect to spine switches. In a spine-leaf architecture, spine switches are used to provide bandwidth and redundancy for leaf switches. Therefore, spine switches do not connect to devices other than leaf switches. As the name implies, spine switches are the backbone of the architecture.
Leaf switches connect to Cisco Application Policy Infrastructure Controllers (APICs). The leaf switch to which a Cisco APIC is directly connected is the first device in a spine-leaf architecture that will be discovered by and registered with the APIC. When a Cisco APIC begins the switch discovery process, it first detects only the leaf switch to which it is connected. After that leaf switch is registered, the APIC discovers each of the spine switches to which the leaf switch is connected. Spine switches do not automatically register with the APIC. When a spine switch is registered with the APIC, the APIC will discover all the leaf switches that are connected to that spine switch.
Leaf switches connect to routers. Routers are typically used to connect to the Internet or to a wide area network (WAN). Leaf switches in a spine-leaf architecture directly connect to routers.
Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 12: ACI Architecture, Spine-Leaf Data Center Design, pp. 447-449

QUESTION 29
Which of the following does not produce API output in JSON format?

A. a REST API
B. A SOAP API
C. GraphQL

D. Falcor

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, only a Simple Object Access Protocol (SOAP) does not produce Application Programming Interface (API) output in JavaScript Object Notation (JSON) format. SOAP is an older API messaging protocol that uses Hypertext Transfer Protocol (HTTP) and Extensible Markup Language (XML) to enable communication between devices. SOAP APIs are typically more resource-intensive than more modern APIs and, therefore, slower. Open APIs can be used to enable services such as billing automation and centralized management of cloud infrastructure.
Representational state transfer (REST) produces output in either JSON or XML format. REST is an API architecture that uses HTTP or HTTP Secure (HTTPS) to enable external resources to access and make use of programmatic methods that are exposed by the API. For example, a web application that retrieves user product reviews from an online marketplace for display on third-party websites might obtain those reviews by using methods provided in an API that is developed and maintained by that marketplace. REST APIs can return data in XML format or in JSON format.
Graph Query Language (GraphQL) produces output in JSON format. GraphQL is an API query language and a runtime that is intended to lower the burden of making multiple API calls to obtain a single set of data. For example, data that requires three or four HTTP GET requests when constructed from a standard REST API might take only one request when using GraphQL. Similar to REST API, GraphQL output is in JSON format. Although GraphQL can use HTTP or HTTPS, it is not limited to those protocols.
Falcor produces output in JSON format. Falcor was developed by Netflix to transport its user interfaces. Similar to GraphQL, Falcor is an attempt to simplify the process of querying an API for remote data by reducing the number of requests that are required to retrieve the data.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 15: Cloud Computing, Application Programming Interfaces, pp. 577-579

QUESTION 30
Which of the following are Type 1 hypervisors? (Choose three.)
A. Microsoft Hyper-V
B. XenServer
C. VMware ESXi
D. Oracle VirtualBox
E. VMware Workstation

Correct Answer: ABC
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available options, Microsoft Hyper-V, VMware ESXi, and XenServer are all Type 1 hypervisors. A hypervisor is software that is capable of virtualizing the physical components of computer hardware. Virtualization enables the creation of multiple virtual machines (VMs) that can be configured and run in separate instances on the same hardware. In this way, virtualization is capable of reducing an organization's expenses on hardware purchases. A Type 1 hypervisor is a hypervisor that is installed on a bare metal server, meaning that the hypervisor is also its own operating system (OS). Because of their proximity to the physical hardware, Type 1 hypervisors typically perform better than Type 2 hypervisors.
VMware Workstation and Oracle VirtualBox are both Type 2 hypervisors. Unlike a Type 1 hypervisor, a Type 2 hypervisor cannot be installed on a bare metal server. Instead, Type 2 hypervisors are applications that are installed on host OSes, such as Microsoft Windows, Mac OS, or Linux. These applications, which are also called hosted hypervisors, use calls to the host OS to translate between guest OSes in VMs and the server hardware. Because they are installed similar to other applications on the host OS, Type 2 hypervisors are typically easier to deploy and maintain than Type 1 hypervisors.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 11: Server Virtualization Solutions, Hypervisor, pp. 407-408

QUESTION 31
You issue the following commands on SwitchA and SwitchB, which are Cisco Nexus 7000 Series switches:

```
SwitchA(config)#vpc domain 101
SwitchA(config-vpc-domain)#peer-keepalive destination 192.168.1.2 source 192.168.1.1 vrf default SwitchA(config-
vpc-domain)#exit
SwitchA(config)#interface range ethernet 2/1 - 2
SwitchA(config-if-range)#switchport
SwitchA(config-if-range)#channel-group 1 mode active
SwitchA(config-if-range)#interface port-channel 1
SwitchA(config-if)#switchport mode trunk
SwitchA(config-if)#vpc peer-link
SwitchB(config)#vpc domain 101
SwitchB(config-vpc-domain)#peer-keepalive destination 192.168.1.1 source 192.168.1.2 vrf default SwitchB(config-
vpc-domain)#exit
SwitchB(config)#interface range ethernet 2/1 - 2
SwitchB(config-if-range)#switchport
SwitchB(config-if-range)#channel-group 1 mode active
SwitchB(config-if-range)#interface port-channel 1
SwitchB(config-if)#switchport mode access
SwitchB(config-if)#vpc peer-link
```

Which of the following is most likely a problem with this configuration?

A. The Ethernet port range is using the wrong channel group mode on SwitchB.
B. The vPC domain ID on SwitchB should not be the same as the value on SwitchA.
C. Port-channel 1 on both switches should be a trunk port.
D. Port-channel 1 on both switches should be an access port.
E. The **vpc peer-link** command should be issued only on SwitchA.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Most likely, the problem with the configuration in this scenario is that Port-channel 1 on both switches should be a trunk port, not just the configuration on SwitchA. Trunk ports are used to carry traffic from multiple virtual local area networks (VLANs) across physical switches. Access ports can only carry data from a single VLAN and are typically connected to end devices, such as hosts or servers. To configure a virtual port channel (vPC) domain between two switches, you should first enable the vPC feature, then configure the vPC domain and peer keepalive links. Next, you should configure 10-gigabit-per-second (Gbps) links between the switches as members of a port channel and configure that port channel as a trunk link. Finally, you should configure the port channel on each switch as a vPC peer link.
The vPC domain ID on SwitchB should be the same as the value on SwitchA. A vPC domain is comprised of two switches per domain. Each switch in the vPC domain must be configured with the same vPC domain ID. To enable vPC configuration on a Cisco Nexus 7000 Series switch, you should issue the **feature vpc** command on both switches. To assign the vPC domain ID, you should issue the **vpc domain** *domain-id* command, where *domain-id* is an integer in the range from 1 through 1000, in global configuration mode. For example, issuing the **vpc domain 101** command on a Cisco Nexus 7000 Series switch configures the switch with a vPC domain ID of 101.
The **vpc peer-link** command should be issued on both switches in this scenario. A vPC peer link should always be comprised of 10-Gbps Ethernet ports. Peer links are configured as a port channel between the two members of the vPC domain. You should configure vPC peer links after you have successfully configured a peer keepalive link. Cisco recommends connecting two 10-Gbps Ethernet ports from two different input/output (I/O) modules. To configure a peer link, you should issue the **vpc peer-link** command in interface configuration mode. For example, the following commands configure a peer link on Port-channel 1:

```
SwitchA(config)#interface port-channel 1
SwitchA(config-if)#switchport mode trunk SwitchA(config-
if)#vpc peer-link
SwitchB(config)#interface port-channel 1
SwitchB(config-if)#switchport mode trunk
SwitchB(config-if)#vpc peer-link
```

It is not a problem that the channel group mode is configured to active on the Ethernet ports in this scenario. It is important to issue the correct **channel-group** commands on a port channel's member ports prior to configuring the port channel. For example, if you are creating Port-channel 1 by using the Ethernet 2/1 and Ethernet 2/2 interfaces, you could issue the following commands on each switch to correctly configure those interfaces as members of the port channel:

```
SwitchA(config)#interface range ethernet 2/1 - 2
```

```
SwitchA(config-if-range)#switchport
SwitchA(config-if-range)#channel-group 1 mode active
SwitchB(config)#interface range ethernet 2/1 - 2
SwitchB(config-if-range)#switchport
SwitchB(config-if-range)#channel-group 1 mode active
```

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 1: Data Center Networking, Configuration Steps of vPC, pp. 27-28

QUESTION 32
Which of the following Cisco UCS servers combine standalone server capabilities with blade server swapability?

A. C-Series servers
B. UCS Mini servers
C. B-Series servers
D. E-Series servers

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available options, Cisco Unified Computing System (UCS) E-Series servers combine standalone server capabilities with blade server swapability. Cisco UCS E-Series servers have similar capabilities to the standalone C-Series servers and do not require connectivity to a UCS fabric. In a small office environment, Cisco UCS E-Series servers are capable of providing the network connectivity and capabilities of a C-Series server along with the availability of B-Series servers.
Cisco UCS C-Series servers are not hot-swappable, nor are they blade servers. Cisco UCS C-Series servers are rack-mount standalone servers that can operate either with or without integration with Cisco UCS Manager. For administrators who are more familiar with traditional Ethernet networks than UCS Fabric Interconnect, C-Series servers will most likely be simpler to deploy and feel more familiar than other Cisco UCS server products.
Cisco UCS B-Series servers are hot-swappable, but they cannot operate as standalone servers. Cisco UCS B-Series servers are blade servers that are installed in a UCS blade chassis. These blade servers can connect only to Cisco UCS Fabric Interconnect, not directly to a traditional Ethernet network. Blade servers in a chassis are typically hot-swappable, unlike the components of a rack-mount server. Therefore, blade server configurations are less likely to result in prolonged downtime if hardware fails.
Cisco UCS Mini servers do not operate as standalone servers. Cisco UCS Mini servers are a compact integration of a Cisco UCS 5108 blade chassis, Cisco UCS 6324 Fabric Interconnects, and Cisco UCS Manager. Unlike other Cisco UCS servers, the UCS Mini server requires a special version of Cisco UCS Manager for management. It is possible to connect Cisco UCS C-Series servers to Cisco UCS Mini servers in order to expand their abilities.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 7: Cisco UCS Architecture, Cisco UCS E-Series Servers, p. 271

QUESTION 33
Which of the following statements about vPC technology limitations is true?

A. vPC forms a Layer 3 port channel.
B. There are only two vPC domain IDs per switch.
C. There is only one switch per vPC domain.
D. Each VDC is a separate switch.
E. vPC peer links are always 1 Gbps.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Each virtual device context (VDC) is a separate switch in a virtual port channel (vPC) domain. A VDC logically virtualizes a switch. A VDC is a single virtual instance of physical switch hardware. A vPC logically combines ports from multiple switches into a single port-channel bundle. Conventional port channels, which are typically used to create high-bandwidth trunk links between two switches, require that all members of the bundle exist on the same switch. vPCs enable virtual domains that are comprised of multiple physical switches to connect as a single entity to a fabric extender, server, or other device.
A vPC domain is comprised of two switches per domain. In addition, a vPC domain cannot be comprised of more than two switches. Each switch in the vPC domain must be configured with the same vPC domain ID. To enable vPC configuration on a Cisco Nexus 7000 Series switch, you should issue the **feature vpc** command on both switches. To assign the vPC domain ID, you should issue the **vpc domain** *domain-id* command, where *domain-id* is an integer in the range from 1 through 1000, in global configuration mode. For example, issuing the **vpc domain 101** command on a Cisco Nexus 7000 Series switch configures the switch with a vPC domain ID of 101.
Only one vPC domain can be configured per switch. If you were to issue more than one **vpc domain** *domain-id* command on a Cisco Nexus 7000 Series switch, the vPC domain ID of the switch would become whatever value was issued last. After you issue the **vpc domain** *domain-id* command, the switch is placed into vPC domain configuration mode. In vPC domain configuration mode, you should configure a peer keepalive link.
Peer keepalive links monitor the remote device to ensure that it is operational. You can configure a peer keepalive link in any virtual routing and forwarding (VRF) instance on the switch. Each switch must use its own Internet Protocol (IP) address as the peer keepalive link source IP address and the remote switch's IP address as the peer keepalive link destination IP address. The following commands configure a peer keepalive link between SwitchA and SwitchB in vPC domain 101:

```
SwitchA(config)#vpc domain 101
SwitchA(config-vpc-domain)#peer-keepalive destination 192.168.1.2 source 192.168.1.1 vrf default
SwitchB(config)#vpc domain 101
```

```
SwitchB(config-vpc-domain)#peer-keepalive destination 192.168.1.1 source 192.168.1.2 vrf default
```

A vPC peer link should always be a 10-gigabit-per-second (Gbps) Ethernet port, not 1 Gbps. Peer links are configured as a port channel between the two members of the vPC domain. You should configure vPC peer links after you have successfully configured a peer keepalive link. Cisco recommends connecting two 10-Gbps Ethernet ports from two different input/output (I/O) modules. To configure a peer link, you should issue the **vpc peer-link** command in interface configuration mode.

```
SwitchA(config)#interface port-channel 1
SwitchA(config-if)#switchport mode trunk SwitchA(config-
if)#vpc peer-link
SwitchB(config)#interface port-channel 1
SwitchB(config-if)#switchport mode trunk
SwitchB(config-if)#vpc peer-link
```

It is important to issue the correct **channel-group** commands on a port channel's member ports prior to configuring the port channel. For example, if you are creating Port-channel 1 by using the Ethernet 2/1 and Ethernet 2/2 interfaces, you could issue the following commands on each switch to correctly configure those interfaces as members of the port channel:

```
SwitchA(config)#interface range ethernet 2/1-2
Switch(config-if-range)#switchport
SwitchA(config-if-range)#channel-group 1 mode active
SwitchB(config)#interface range ethernet 2/1-2
SwitchB(config-if-range)#switchport
SwitchB(config-if-range)#channel-group 1 mode active
```


A vPC forms an Open System Interconnection (OSI) networking model Layer 2 port channel, not a Layer 3 port channel. The vPC feature is not capable of supporting Layer 3 port channels. Therefore, any routing that is configured from the vPC peers to other parts of the network should be performed on separate Layer 3 ports.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 1: Data Center Networking, vPC Limitations, p. 27

QUESTION 34

Which of the following Cisco UCS servers is a rack-mount server that does not require a UCS fabric?

A. E-Series servers
B. UCS Mini servers
C. B-Series servers
D. C-Series servers

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available options, Cisco Unified Computing System (UCS) C-Series servers do not require a UCS fabric. Cisco UCS C-Series servers are rack-mount standalone servers that can operate either with or without integration with Cisco UCS Manager. For administrators who are more familiar with traditional Ethernet networks than UCS Fabric Interconnect, C-Series servers will most likely be simpler to deploy and feel more familiar than other Cisco UCS server products.
Cisco UCS B-Series servers require a UCS fabric. Cisco UCS B-Series servers are blade servers that are installed in a UCS blade chassis. These blade servers can connect only to Cisco UCS Fabric Interconnect, not directly to a traditional Ethernet network. Blade servers in a chassis are typically hot-swappable, unlike the components of a rack-mount server. Therefore, blade server configurations are less likely to result in prolonged downtime if hardware fails.
Cisco UCS E-Series servers are blade servers. However, Cisco UCS E-Series servers have similar capabilities to the standalone C-Series servers and do not require connectivity to a UCS fabric. In a small office environment, Cisco UCS E-Series servers are capable of providing the network connectivity and capabilities of a C-Series server along with the availability of B-Series servers.
Cisco UCS Mini servers are a compact integration of a Cisco UCS 5108 blade chassis, Cisco UCS 6324 Fabric Interconnects, and Cisco UCS Manager. Unlike other Cisco UCS servers, the UCS Mini server requires a special version of Cisco UCS Manager for management. It is possible to connect Cisco UCS C-Series servers to Cisco UCS Mini servers in order to expand their abilities.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 7: Cisco UCS Architecture, Cisco UCS C-Series Rackmount Servers, pp. 258-261

QUESTION 35
Which of the following Cisco UCS servers typically reside in a blade chassis? (Choose three.)

A. E-Series servers
B. UCS Mini servers
C. M-Series servers
D. B-Series servers
E. C-Series servers

Correct Answer: ABD
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, Cisco Unified Computing System (UCS) B-Series servers, Cisco UCS E-Series servers, and Cisco UCS Mini servers typically reside in a blade chassis. Cisco UCS B-Series servers are blade servers that are installed in a UCS blade chassis. These blade servers can connect only to Cisco UCS

Fabric Interconnect, not directly to a traditional Ethernet network. Blade servers in a chassis are typically hot-swappable, unlike the components of a rack-mount server. Therefore, blade server configurations are less likely to result in prolonged downtime if hardware fails.
Cisco UCS E-Series servers are typically installed in a blade chassis. However, Cisco UCS E-Series servers have similar capabilities to the standalone C-Series servers and do not require connectivity to a UCS fabric. In a small office environment, Cisco UCS E-Series servers are capable of providing the network connectivity and capabilities of a C-Series server along with the availability of B-Series servers.
Cisco UCS Mini servers are a compact integration of a Cisco UCS 5108 blade chassis, Cisco UCS 6324 Fabric Interconnects, and Cisco UCS Manager. Unlike other Cisco UCS servers, the UCS Mini server requires a special version of Cisco UCS Manager for management. It is possible to connect Cisco UCS C-Series servers to Cisco UCS Mini servers in order to expand their abilities.
Cisco UCS C-Series servers are typically installed directly in a rack. Cisco UCS C-Series servers are rack-mount standalone servers that can operate either with or without integration with Cisco UCS Manager. Therefore, Cisco UCS C-Series servers do not require a UCS fabric. For administrators who are more familiar with traditional Ethernet networks than UCS Fabric Interconnect, C-Series servers will most likely be simpler to deploy and feel more familiar than other Cisco UCS server products.
Cisco UCS M-Series servers are neither rack servers nor blade servers. M-Series servers are comprised of computing modules that are inserted into a modular chassis. The chassis, not the server, can be installed in a rack. The modular design enables the segregation of computing components from infrastructure components.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 7: Cisco UCS Architecture, Cisco UCS B-Series Blade Servers, pp. 251-253
Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 7: Cisco UCS Architecture, Cisco UCS E-Series Servers, p. 271
Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 7: Cisco UCS Architecture, Cisco UCS Mini, p. 272
Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 7: Cisco UCS Architecture, Cisco UCS M-Series Servers, pp. 272-273

QUESTION 36
You are configuring a service profile for a Cisco UCS server that contains four physical NICs but no physical HBAs.
How many vHBAs can be configured for this server?

A. none
B. eight
C. two
D. four

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

No virtual host bus adapters (vHBAs) can be configured for a Cisco Unified Computing System (UCS) server that contains four physical network interface cards (NICs) but no physical host bus adapters (HBAs). A Cisco UCS can be configured with the number of vHBAs that corresponds to available physical HBAs on the physical adapters that are installed in the device. Similarly, a Cisco UCS server can be configured with the number of virtual NICs (vNICs) that corresponds to available physical NICs on the server.
It is not possible to configure two vHBAs for the Cisco UCS server in this scenario. In order to configure two vHBAs, the server would need to be configured with one or more adapters that each contains one or more physical HBAs. For example, a Cisco converged network adapter typically contains two physical ports. Therefore, a Cisco UCS server that is configured with a single converged network adapter could be configured with two vHBAs. A converged network adapter is a single unit that combines a physical HBA and a physical Ethernet NIC. A Cisco converged network adapter typically contains two of these types of ports. It is not possible to configure four vHBAs for the Cisco UCS server in this scenario. To configure four vHBAs, the UCS server would need to be configured with adapters that contain up to four physical HBAs.
It is not possible to configure eight vHBAs for the Cisco UCS server in this scenario. To configure eight vHBAs, the UCS server would need to be configured with adapters that contain up to eight physical HBAs.

Reference:

Cisco: Overview of Cisco Unified Computing System: Configuration through Service Profiles

QUESTION 37
Which of the following FIP protocols finds forwarders that can accept logins?

A. FDISC
B. FIP VLAN Discovery
C. FIP FCE Discovery
D. FLOGI

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, only FCoE Initialization Protocol (FIP) Fibre Channel Forwarder (FCF) Discovery finds forwarders that can accept logins. FIP FCF Discovery is one of two discovery protocols used by FIP during the Fibre Channel over Ethernet (FCoE) initialization process. FIP itself is a control protocol that is used to create and maintain links between FCoE device pairs, such as FCoE nodes (ENodes) and FCFs. ENodes are FCoE entities that are similar to host bus adapters (HBAs) in native Fibre Channel (FC) networks. FCFs, on the other hand, are FCoE entities that are similar to FC switches in native FC networks.
FIP FCF Discovery is typically the second discovery process to occur during FCoE initialization. This process enables ENodes to discover FCFs that allow logins. FCFs that allow logins send periodic FCF Discovery advertisements on each FCoE virtual local area network (VLAN).

FIP VLAN Discovery is typically the first discovery process to occur during FCoE initialization. This process discovers the VLAN that should be used to send all other FIP traffic during the initialization. This same VLAN is also used by FCoE encapsulation. The FIP VLAN Discovery protocol is the only FIP protocol that runs on the native VLAN.
Fabric Login (FLOGI) and Fabric Discovery (FDISC) messages comprise the final protocol in the FCoE FIP initialization process. These messages are used to activate Open Systems Interconnection (OSI) network model Layer 2, or the Data link layer, of the fabric. It is at this point in the initialization process that an FC ID is assigned to the N port, which is the port that connects the node to the switch.

Reference:
Cisco: FCoE Initiation Protocol

QUESTION 38
Which of the following are Type 2 hypervisors? (Choose three.)

A. VMware Workstation
B. VMware ESXi
C. Microsoft Hyper-V
D. VMWare Fusion
E. Oracle VirtualBox

Correct Answer: ADE
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available options, VMware Fusion, VMware Workstation, and Oracle VirtualBox are all Type 2 hypervisors. A hypervisor is software that is capable of virtualizing the physical components of computer hardware. Virtualization enables the creation of multiple virtual machines (VMs) that can be configured and run in separate instances on the same hardware. In this way, virtualization is capable of reducing an organization's expenses on hardware purchases.
Type 2 hypervisors are applications that are installed on host operating systems (OSs), such as Microsoft Windows, Mac OS, or Linux. These applications, which are also called hosted hypervisors, use calls to the host OS to translate between guest OSs in VMs and the server hardware. Because they are installed similar to other applications on the host OS, Type 2 hypervisors are typically easier to deploy and maintain than Type 1 hypervisors.
Microsoft Hyper-V and VMware ESXi are both Type 1 hypervisors. A Type 1 hypervisor is a hypervisor that is installed on a bare metal server, meaning that the hypervisor is also its own OS. Unlike a Type 1 hypervisor, a Type 2 hypervisor cannot be installed on a bare metal server. Because of their proximity to the physical hardware, Type 1 hypervisors typically perform better than Type 2 hypervisors.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 11: Server Virtualization Solutions, Hypervisor, pp. 407-408

QUESTION 39
Which of the following logically virtualizes a Cisco Nexus 7000 Series switch?

A. a VRF instance
B. a VIC
C. a VDC
D. a vPC
Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

A virtual device context (VDC) logically virtualizes a Cisco Nexus 7000 Series switch. A VDC is a single virtual instance of physical switch hardware. By default, the control plane of the Cisco Nexus 7000 Series switch is configured to run a single VDC. It is possible to configure multiple VDCs on the same hardware. A single VDC can contain multiple virtual local area networks (VLANs) and virtual routing and forwarding (VRF) instances.
A virtual port channel (vPC) does not logically virtualize a Cisco Nexus 7000 Series switch. A vPC enables ports from multiple switches to be combined into a single port channel bundle. Conventional port channels, which are typically used to create high-bandwidth trunk links between two switches, require that all members of the bundle exist on the same switch. vPCs enable virtual domains that are comprised of multiple physical switches to connect as a single entity to a fabric extender, server, or other device.
A VRF instance does not logically virtualize a Cisco Nexus 7000 Series switch. A VRF is a virtual instance of an Open Systems Interconnection (OSI) network model Layer 3 address domain. VRFs enable a router to maintain multiple, simultaneous routing tables. Therefore, VRFs can be configured on a single router to serve multiple Layer 3 domains instead of implementing multiple hardware routers.

A virtual interface card (VIC) does not logically virtualize a Cisco Nexus 7000 Series switch. A VIC is a Cisco device that can be used to create multiple logical network interface cards (NICs) and host bus adapters (HBAs). VICs such as the Cisco M81KR send Fibre Channel over Ethernet (FCoE) traffic and normal Ethernet traffic over the same physical medium.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 6: Virtualizing Cisco Network Devices, Describing VDCs on the Cisco Nexus 7000 Series Switch, pp. 192-194

QUESTION 40
Which of the following typically defines a Layer 2 address space and flood domain within a Cisco ACI fabric?

A. an application profile
B. a VRF instance

C. a bridge domainD. an EPG

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Of the available choices, a bridge domain typically defines a Layer 2 address space and flood domain within a Cisco Application Centric Infrastructure (ACI) fabric. An ACI bridge domain is similar to a traditional networking virtual local area network (VLAN) in that it is an Open Systems Interconnection (OST) networking model Layer 2 broadcast domain. Bridge domains define the Media Access Control (MAC) address space. Typically, a bridge domain is associated with a single subnet, although multiple subnets can be associated within a given bridge domain. Bridge domains are typically connected to virtual routing and forwarding (VRF) instances within a given ACI tenant.
A VRF instance is a Layer 3 forwarding domain within a given ACI tenant. VRF instances are also known as contexts or private networks. Multiple bridge domains can be connected to a given VRF instance within a tenant.
An application profile is used to configure policies and relationships between endpoint groups (EPGs). Typically, an application profile is a container for one or more logically related EPGs. For example, EPGs that provide similar services or functions might be associated with the same application profile.
An EPG is a logical ACI construct that contains multiple related endpoints. Endpoints are physical or virtual devices that are connected to a network. For example, a web server is an endpoint. EPGs contain endpoints that have similar policy requirements, although not necessarily similar functions. EPGs enable group management of the policies for the endpoints they contain.

Reference:

Cisco: Cisco Application Centric Infrastructure Fundamentals: Bridge Domains and Subnets

QUESTION 41
Which of the following devices integrates with a Cisco Nexus 7000 Series switch and is used to automatically scale up a Layer 2 topology?

A. the Cisco Nexus 9000 Series switch
B. the Cisco Nexus 1000v switch
C. the Cisco Nexus 2000 Series switch
D. the Cisco Nexus 5000 Series switch

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, only the Cisco Nexus 2000 Series switch integrates with a Cisco Nexus 7000 Series switch and is used to automatically scale up an Open Systems Interconnection (OSI) networking model Layer 2 topology. The Cisco Nexus 2000 Series of switches are fabric extenders (FEXs) and cannot operate as standalone switches. By default, Cisco FEX devices are not configured with software. Instead, they depend on parent switches, such as a Cisco Nexus 7000 Series switch, to install and configure software, provide forwarding tables, and provide control plane functionality.
The Cisco Nexus 1000v is a virtual switch that is capable of connecting to upstream physical switches in order to provide connectivity for a virtual machine (VM) network environment. Although the Cisco Nexus 1000v operates similar to a standard switch, it exists only as software in a virtual environment and is therefore not a physical switch.
Cisco Nexus 5000 Series switches operate as standalone physical switches. Cisco Nexus 5000 Series switches are data center access layer switches that can support 10-gigabit-per-second (Gbps) or 40-Gbps Ethernet, depending on the model. Native Fibre Channel (FC) and FC over Ethernet (FCoE) are also supported by Cisco Nexus 5000 Series switches.
Cisco Nexus 9000 Series switches operate as standalone physical switches. Cisco Nexus 9000 Series switches can operate either as traditional NX-OS switches or in an Application Centric Infrastructure (ACI) mode. Unlike Cisco Nexus 7000 Series switches, Cisco Nexus 9000 Series switches do not support virtual device contexts (VDCs) or storage protocols.

Reference:

Cisco: Cisco Nexus 2000 Series NX-OS Fabric Extender Software Configuration Guide for Cisco Nexus 7000 Series Switches, Release 6.x: Information About the Cisco Nexus 2000 Series Fabric Extender

QUESTION 42
You are creating a workflow in UCS Director's Workflow Designer. You have connected each task's On Success event in the workflow to the appropriate next task. You now need to connect each task's On Failure event to the appropriate task. Which of the following tasks are you most likely to choose?

A. completed (success)
B. completed (failed)
C. start
D. the next task in the workflow

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, you are most likely to connect each task’s On Failure event to the completed (failed) task. Workflows determine the order in which tasks that are designed to automate complex IT operations are performed. Workflow Designer allows administrators to create workflows that can then be automated by using Unified Computing System (UCS) Director's orchestrator.

The following Cisco UCS Director's Workflow Designer tasks are predefined when a workflow is created: ▪
Completed (failed)
▪ Completed (success)
▪ Start
The start task is the beginning of the workflow. The completed (failed) task represents the end of a workflow when the desired result could not be achieved. The completed (success) task represents a successfully completed workflow. Each task in a workflow processes input and produces output that is sent to the next task in the workflow. In addition, each task contains an On Success event and an On Failure event that can be used to determine which task should be performed next based on whether the task could be successfully completed. On Success events should be connected to the next task in the workflow. On Failure events, on the other hand, should be connected to the completed (failed) task so that the workflow does not attempt to perform more tasks that would rely on successful output from the previously failed task.
Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 17: Understanding and Troubleshooting UCSD Workflows, Creating Workflows, pp. 641-645

QUESTION 43
Which of the following is Cisco software that can be used to construct a private cloud?

A. Cisco UCS Central
B. Cisco UCS Manager
C. Cisco IMC Supervisor
D. Cisco UCS Director

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, only Cisco Unified Computing System (UCS) Director is software that can automate actions and be used to construct a private cloud. Cisco UCS Director creates a basic Infrastructure as a Service (IaaS) framework by using hardware abstraction to convert hardware and software into programmable actions that can then be combined into an automated custom workflow. Thus Cisco UCS Director enables administrators to construct a private cloud in which they can automate and orchestrate both physical and virtual components of a data center. Cisco UCS Director is typically accessed by using a web-based interface.
Cisco UCS Central cannot be used to construct a private cloud. Cisco UCS Central is software that can be used to manage multiple UCS domains, including domains that are separated by geographical boundaries. Cisco UCS Central can be used to deploy standardized configurations and policies from a central virtual machine (VM).
Cisco UCS Manager cannot be used to construct a private cloud. Cisco UCS Manager is web-based software that can be used to manage a single UCS domain.

The software is typically embedded in Cisco UCS fabric interconnects rather than installed in a VM or on separate physical servers.
Cisco Integrated Management Controller (IMC) Supervisor cannot be used to construct a private cloud. Cisco IMC Supervisor is software that can be used to centrally manage multiple standalone Cisco C-Series and E-Series servers. The servers need not be located at the same site. Cisco IMC Supervisor uses a webbased interface and is typically deployed as a downloadable virtual application.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 16: UCS Director, What Is UCS Director?, pp. 587-588.
Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 7: Cisco UCS Architecture, Cisco UCS Software, pp. 274-277

QUESTION 44
You want to allow remote users to log in to a Nexus 7000 Series switch nondefault VDC by using TACACS+. The TACACS+ configuration has been previously completed on the switch. You issue the following commands:
switchto vdc MyVDC
configure terminal
aaa user default-role
exit copy running-config start-
config
Which of the following user roles will occur when a remote user logs in to the VDC named MyVDC by using TACACS+?

A. The user will be assigned the vdc-operator role.
B. The user will be assigned the network-admin role.
C. The user will not be assigned a role and will be denied login.
D. The user will be assigned the vdc-admin role.
E. © The user will be assigned the network-operator role.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

The user will be assigned the vdc-operator role when the remote user logs in by using Terminal Access Controller Access-Control System Plus (TACACS+) in this scenario. The vdc-operator role has read-only access to a specific virtual device context (VDC) on the switch. In this scenario, the **aaa user default-role** command has been issued in the VDC named MyVDC, which is a nondefault VDC on the switch. The **aaa user default-role** command configures the Authentication, Authorization, and Accounting (AAA) feature on the switch to automatically assign remote users the default user role at login. The default remote user role for nondefault VDCs on a Cisco Nexus switch is the vdc-operator role.

Cisco Nexus switches use role-based access control (RBAC) to assign management privileges to a given user. By default, a Nexus 7000 switch is configured with the following user roles:
▪ network-admin — has read and write access to all VDCs on the switch ▪ network-operator — has read-only access to all the VDCs on the switch ▪ vdc-admin — has read and write access to a specific VDC on the switch ▪ vdc-operator — has read-only access to a specific VDC on the switch
The user will not be assigned the network-admin role. In addition, the user will not be assigned the network-operator role. These roles are applied to users who have access to all VDCs that are configured on the switch, not a specific nondefault VDC. If the **aaa user default-role** command had been issued in the default VDC in this scenario, remote users who log in to the default VDC would be assigned a network-operator user role.
The user will not be assigned the vdc-admin user role. The vdc-admin user role allows read and write access to a specific VDC on the switch. If remote users were automatically assigned the vdc-admin role when logging in to the VDC named MyVDC, those users would have administrative access to the VDC, which is a security risk.
The user will be assigned a role and will not be denied login. In this scenario, TACACS+ is already configured on the Cisco Nexus 7000 Series switch. In addition, the **aaa user default-role** command has been issued. If the command had not been issued or if the **no aaa user default-role** command had been issued later in the configuration, remote users who attempt to log in to the VDC named MyVDC would be denied access because no user role is assigned to those users.

Reference:

Cisco: Configuring AAA: Enabling the Default User Role for AAA Authentication

QUESTION 45
You issue the **port-channel load-balance scr-dst-port module 4** command on a Cisco Nexus 7000 Series switch running NX-OS 6.2.
How is distribution loaded for port channels on slot 4?

A. based on the source and destination port
B. based on the source and destination MAC address
C. based on the defaults because the command contains invalid syntax
D. based on the source IP address and destination port
E. based on the source MAC address and destination port only

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Distribution for port channels on slot 4 in this scenario is loaded based on the source and destination port because the command has been issued with the **src-dstport** keyword and the **module 4** keyword. The **port-channel load-balance** command can be used to modify the load distribution criteria used by EtherChannel. The basic syntax of the **port-channel load-balance** command is **port-channel load-balance** *method* [**module** *slot*], where *method* is one of the following 10 keywords: ▪ **dst-ip** ▪ **dst-mac** ▪ **dst-port** ▪ **src-dst-ip** ▪ **src-dst-mac** ▪ **src-dst-port** ▪ **src-ip** ▪ **src-mac** ▪ **src-port** ▪ **vlan-only**
The **dst-ip** keyword, **dst-mac** keyword, and **dst-port** keyword load port channel distribution based on the destination Internet Protocol (IP) address, Media Access Control (MAC) address, and port number, respectively. Similarly, the **src-ip** keyword, **src-mac** keyword, and **src-port** keyword load port channel distribution based on the source IP address, source MAC address, and port number, respectively. The **src-dst-ip** keyword, **src-dst-mac** keyword, and **src-dst-port** keyword load port channel distribution based on the source and destination IP addresses, MAC addresses, and port numbers, respectively. The **vlan-only** keyword loads distribution on only the virtual local area network (VLAN) modules.
The optional **module** keyword accepts a slot number value. If you configure the **port-channel load-balance** command with the **module** keyword, the configuration applies only to the specified slot. Otherwise, the configuration applies to the entire device. By default on a Cisco Nexus switch, a port channel load balances Layer 2 packets based on the source and destination MAC addresses. Layer 3 packets, on the other hand, are load balanced based on the source and destination IP addresses. You must be operating in the default virtual device context (VDC) on the switch in order to issue the **port-channel load-balance** command.
The command in this scenario does not contain invalid syntax, because the Nexus switch in this scenario is running NX-OS 6.2. Prior to NX-OS 5.1(3), the **portchannel load-balance** command required an **ethernet** keyword. Therefore, the valid syntax for NX-OS versions older than 5.1(3) is **port-channel load-balance ethernet** *method* [**module** *slot*]. In NX-OS 5.1(3), Cisco removed that keyword.

Reference:

Cisco: Cisco Nexus 7000 Series NX-OS Interfaces Command Reference: port-channel load-balance

QUESTION 46
You want to migrate a VMware VM and its datastore from one ESXi server to another ESXi server. You do not want to power off the VM to perform the migration. Which of the following solutions should you choose?

A. copying or cloning
B. cold migration
C. Storage vMotion
D. vSphere vMotion

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, you should choose Storage vMotion if you want to migrate a VMware virtual machine (VM) and its datastore from one VMware ESXi server to another ESXi server without powering off the VM. VMware's ESXi server is a bare-metal server virtualization technology, which means that ESXi is installed directly on the hardware it is virtualizing instead of running on top of another operating system (OS). This layer of hardware abstraction enables tools like Storage vMotion to migrate ESXi VMs and their datastores from one host to another without powering off the VM, enabling the VM's users to continue working without interruption.
You should not choose vSphere vMotion to perform the migration in this scenario, because vMotion allows migration of only the VM and its virtual components; it does not allow migration of the datastore. The datastore is the repository of VM-related files, such as logs and virtual disks. When migrating a VM by using vMotion, only the virtualized environment moves to a new host, not the datastore.
You should not choose cold migration, copying, or cloning in this scenario. Cold migration is the process of powering down a VM and moving the VM or the VM and its datastore to a new location. While a cold migration is in progress, no users can perform tasks inside the VM. Both copying and cloning create new instances of a given VM. Therefore, neither action is a form of migrating a VM to another host. Typically, a VM must be powered off or suspended in order to successfully copy or clone it.

Reference:

VMware: VMware Docs: Migrating Virtual Machines
VMware: Virtualization Overview

QUESTION 47
Which of the following is a software environment that runs a separate OS from its host OS?

A. a hypervisor
B. a mezzanine card
C. a VM
D. an API

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

A virtual machine (VM) is a software environment that runs a separate operating system (OS) from its host OS. However, a VM does not itself abstract the hardware on which the host OS is installed. VMs provide the hardware environment for guest OSes and the applications they run by making calls to the hypervisor, which then makes calls to either the host OS or the bare-metal server, depending on the type of hypervisor installed on the device.
A hypervisor is software that has two roles: the abstraction of physical hardware and the creation of VMs. Hypervisors are capable of virtualizing the physical components of computer hardware. Virtualization enables the creation of multiple VMs that can be configured and run in separate instances on the same hardware. Hardware abstraction is the use of software to emulate physical hardware. Hardware abstraction enables device-independent software development and allows a given VM to become portable between physical devices.
An Application Programming Interface (API) is typically used to enable an application to perform functions on a remote framework, database, or application. For example, representational state transfer (REST) is an API architecture that uses Hypertext Transfer Protocol (HTTP) or HTTP Secure (HTTPS) to enable external resources to access and make use of programmatic methods that are exposed by the API. A web application that retrieves user product reviews from an online marketplace for display on third-party websites might obtain those reviews by using methods provided in an API that is developed and maintained by that marketplace.
A mezzanine card is a computer hardware component that can be plugged into expansion slots on a main board. Cisco Unified Computing System (UCS) B-Series blade servers use mezzanine cards to add a variety of network interfaces to the system. For example, the Cisco UCS Virtual Interface Card (VIC) 1280 is a mezzanine card that adds 10-gigabit-per-second (Gbps) Ethernet port and Fibre Channel over Ethernet (FCoE) capabilities to a Cisco UCS B-Series blade server.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 11: Server Virtualization Solutions, Virtual Machines, pp. 408-409

QUESTION 48
Which of the following are most likely to operate in the control plane of a Nexus switch? (Choose three.)

A. OSPF
B. cut-through switching
C. BGP
D. store-and-forward switching
E. EIGRP
F. SNMP

Correct Answer: ACE
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, Enhanced Interior Gateway Routing Protocol (EIGRP), Open Shortest Path First (OSPF), and Border Gateway Protocol (BGP) are all most likely to operate in the control plane of a Nexus switch. A Nexus switch consists of three operational planes: the data plane, which is also known as the

forwarding plane, the control plane, and the management plane. The control plane is responsible for gathering and calculating the information required to make the decisions that the data plane needs for forwarding. Routing protocols operate in the control plane because they enable the collection and transfer of routing information between neighbors. This information is used to construct routing tables that the data plane can then use for forwarding.
Cut-through switching and store-and-forward switching are most likely to operate in the data plane of a Nexus switch. Of the three, the data plane is where traffic forwarding occurs. Cut-through switching allows a switch to begin forwarding a frame before the frame has been received in its entirety. Store-and-forward switching receives an entire frame and stores it in memory before forwarding the frame to its destination.
Simple Network Management Protocol (SNMP) is an Internet Protocol (IP) network management protocol that operates in the management plane of a Nexus switch. The management plane is responsible for monitoring and configuration of the control plane. Therefore, network administrators typically interact directly with protocols running in the management plane.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 2: Management and Monitoring of Cisco Nexus Devices, Control Plane, pp. 53-58

QUESTION 49

Which of the following frame fields does Cisco FabricPath use to identify the unique FabricPath topology that a unicast frame is traversing?

A. FTAG
B. IS-IS
C. LID
D. STP
E. TTL

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, Cisco FabricPath uses the forwarding tag (FTAG) frame field to identify the unique FabricPath topology that a unicast frame is traversing. Each topology in FabricPath is assigned a unique tag. For multicast or broadcast traffic, the 10-bit FTAG field contains an ID for a forwarding tree that contains multiple destinations within the topology. The FTAG field is the second of three fields that reside in the FabricPath tag field. It is preceded by the 16-bit Ethertype field and succeeded by the Time to Live (TTL) field.
Cisco FabricPath frames are classic Ethernet frames that are encapsulated with a 16-byte FabricPath header. This header contains a 48-bit outer destination address (ODA), which is a Media Access Control (MAC) address, and a 48-bit outer source address (OSA). In addition, the field contains the 32-bit FabricPath tag. The classic Ethernet frame's cyclic redundancy check (CRC) field is replaced by a new CRC field that is updated to reflect the additional header data in the frame.
Cisco FabricPath does not use the TTL frame field to identify the unique FabricPath topology that a unicast frame is traversing. The TTL field is the Cisco

FabricPath frame field that FabricPath uses to mitigate temporary Open Systems Interconnection (OSI) networking model Layer 2 loops. The TTL field in a FabricPath frame operates similarly to the TTL field in IP networking in that the field is decremented by a value of 1 each time it traverses a new hop. If the TTL expires, the frame is discarded. The TTL field is a 6-bit field that resides at the end of the FabricPath tag field.
Cisco FabricPath does not use an Intermediate System-to-Intermediate System (IS-IS) frame field to identify the unique FabricPath topology that a unicast frame is traversing. In addition, FabricPath does not use a Spanning Tree Protocol (STP) frame field. However, the IS-IS routing protocol is used as a Layer 3 replacement for traditional STP in a Cisco FabricPath topology. In traditional networking, STP is used to prevent Layer 2 switching loops in a topology that contains redundant links. The use of the IS-IS routing protocol ensures that Cisco FabricPath operates as a multipath environment for Layer 2 packets. In other words, IS-IS ensures that Cisco FabricPath is capable of Layer 2 multipath forwarding.
Cisco FabricPath does not use the local ID (LID) frame field to identify the unique FabricPath topology that a unicast frame is traversing. Instead, the LID field stores a 16-bit value that identifies the edge port that a packet is either destined to or sent from. The LID field is the last field in the ODA and OSA fields of a Cisco FabricPath header. The edge port can be either a physical port or a logical port. In addition, the value in the LID field is locally significant to the switch that the frame is traversing.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 1: Data Center Networking, FabricPath Frame Format, pp. 34-35

QUESTION 50

Which of the following Cisco switches can be configured only by first using Telnet or SSH to access a parent device?

A. Nexus 7000 Series
B. Nexus 2000 Series
C. Nexus 5000 Series
D. Nexus 9000 Series

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, Cisco Nexus 2000 Series switches can be configured only by first using Telnet or Secure Shell (SSH) to access a parent device. The Cisco Nexus 2000 Series of switches are fabric extenders (FEXs) and cannot operate as standalone switches. FEX technologies depend on parent switches, such as a Cisco Nexus 5500 Series switch or a Cisco Nexus 7000 Series switch, to provide forwarding tables and control plane functionality. FEX technologies are intended to extend the network to edge devices. Typically, FEX devices in the Cisco Nexus 2000 Series are managed by first connecting to the parent device by using either Telnet or SSH and then configuring the FEX.
Cisco Nexus 5000 Series switches operate as standalone physical switches. Cisco Nexus 5000 Series switches are data center access layer switches that can support 10-gigabit-per-second (Gbps) or 40-Gbps Ethernet, depending on the model. Native Fibre Channel (FC) and FC over Ethernet (FCoE) are also supported by Cisco Nexus 5000 Series switches.

Cisco Nexus 7000 Series switches operate as standalone physical switches. Cisco Nexus 7000 Series switches are typically used as an end-to-end data center solution, which means that the series is capable of supporting all three layers of the data center architecture: core layer, aggregation layer, and access layer. In addition, the Cisco Nexus 7000 Series supports virtual device contexts (VDCs). The Cisco Nexus 7000 Series can support up to 100-Gbps Ethernet.
Cisco Nexus 9000 Series switches operate as standalone physical switches. Cisco Nexus 9000 Series switches can operate either as traditional NX-OS switches or in an Application Centric Infrastructure (ACI) mode. Unlike Cisco Nexus 7000 Series, Cisco Nexus 9000 Series switches do not support VDCs or storage protocols.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 6: Virtualizing Cisco Network Devices, Nexus 2000 Series Fabric Extender Connectivity, pp. 211-212

QUESTION 51
Which of the following statements best describes vPC domains?

A. There can be only two peers per domain.
B. They monitor the status of vPC peers.
C. They synchronize the state between two vPC peers.
D. They synchronize the control plane and the data plane.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

There can be only two peers, or switches, per virtual port channel (vPC) domain. A vPC enables you to bundle ports from two peers, which form a domain, into a single Open Systems Interconnection (OSI) Layer 2 port channel. Similar to a normal port channel, a vPC bundles multiple switch ports into a single high-speed trunk port. A single vPC domain cannot contain ports from more than two switches. For ports on two switches to successfully form a vPC domain, all the following must be true:
- The vPC feature must be enabled on both switches.
- The vPC domain ID must be the same on both switches.
- The peer keepalive link must be configured and must be 10 gigabits per second (Gbps) or more. ▪

The vPC number must be the same on both switches.
A vPC peer link, not a vPC domain, synchronizes the state between two vPC peers. A vPC peer link is typically comprised of a port channel made up of two physical ports on each switch. This link synchronizes Media Access Control (MAC) address tables between switches and serves as a transport for data plane traffic. Bridge protocol data unit (BPDU) and Link Aggregation Control Protocol (LACP) packets are also forwarded to the second peer over this link, which causes the vPC peers to appear to be a single control plane.

A vPC peer keepalive link, not a vPC domain, monitors the status of vPC peers. The peer keepalive link operates at Layer 3 of the OSI networking model; it is used to ensure that vPC switches are capable of determining whether a vPC domain peer has failed. Peer keepalive links can be configured to operate in any virtual routing and forwarding (VRF) instance, including the management VRF. Each vPC peer keepalive link is configured with the remote peer's IP address as its destination IP address and the local peer's IP address as its source address. Peer keepalive links must be trunk links.
Cisco Fabric Services, not a vPC domain, synchronizes the control plane and the data plane. Cisco Fabric Services is a messaging protocol that operates between vPC peers. Control plane and data plane information is synchronized over the vPC peer link.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 1: Data Center Networking, Components of vPC, pp. 22-24

QUESTION 52
An administrator configures a Cisco GSS and migrates DNS services to the GSS.
Which of the following is most likely being implemented?

A. a GSLB service
B. an APIC solution
C. a hypervisor running DNS services
D. a FabricPath architecture

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Most likely, a Cisco global server load balancing (GSLB) service is being implemented if an administrator configures a Cisco Global Site Selector (GSS) and migrates Domain Name System (DNS) services to the GSS. A Cisco GSLB solution is designed to optimize DNS infrastructure, thereby ensuring business continuity in the event of disaster. When DNS services are migrated to GSS, disaster recovery is enhanced by the global load balancing of server load balancers (SLBs) across data centers in disparate geographic locations.
It is not likely that the administrator is implementing a Cisco FabricPath architecture. In addition, it is not likely that the administrator is implementing a Cisco Application Policy Infrastructure Controller (APIC) solution. The Cisco APIC is a means of managing the Cisco Application Centric Infrastructure (ACI). A Cisco ACI architecture requires both the APIC and the spine switches and leaf switches of FabricPath to complete the architecture. The APIC communicates with the spine and leaf nodes and provides policy distribution as well as centralized management.
Spine switches are the Cisco FabricPath component that form the backbone of Cisco FabricPath's switching fabric. Typically, leaf switches are connected to every spine switch along the backbone so that the spine switches provide connectivity between the leaf switches. A leaf switch is the Cisco FabricPath component that provides access layer connectivity. End hosts and classic Ethernet (CE) networks are typically directly connected to leaf switches by using edge ports. Leaf switches connect to spine switches by using core ports.

It is not likely that the administrator is implementing a hypervisor running DNS services. A hypervisor is hardware virtualization software that runs either on a baremetal server or as an application on an operating system (OS); hypervisors are used to create and run one or more virtual machines (VMs). Bare-metal server hypervisors are known as Type 1 hypervisors. Application hypervisors are known as Type 2 hypervisors. Neither type of hypervisor will run DNS services by itself. Instead, a VM would need to be created and configured with a guest OS. The DNS service would then need to be configured within the guest OS.

Reference:

Cisco: Introducing the Global Site Selector: GSLB Using the GSS

QUESTION 53
Which of the following FCoE switch port types might require you to consider an STP configuration?

A. a VF port
B. a SPAN port
C. a VE port
D. a VN port

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

A Fibre Channel over Ethernet (FCoE) virtual fabric (VF) interface port type might require you to consider a Spanning Tree Protocol (STP) configuration. FCoE is used in data centers to encapsulate Fibre Channel (FC) over an Ethernet network. This encapsulation enables the FC protocol to communicate over 10 gigabit-persecond (Gbps) Ethernet. There are two types of FCoE switch ports: a VF port and a virtual edge (VE) port.
An FCoE VE port typically connects to an end host. If the end host is connected to an Ethernet network that is configured with virtual local area networks (VLANs), the STP configuration might require extra attention, especially if the Ethernet fabric is not using Per-VLAN Spanning Tree Plus (PVST+). A proper STP configuration on the Ethernet fabric prevents the Ethernet topology from affecting storage area network (SAN) traffic.
An FCoE VE port typically connects to a port on another FC forwarder (FCF). STP does not operate on VE ports, because these ports typically connect two FCFs. FC does not require switching loop prevention, because FCFs have no concept of switching loops. VE ports typically default to trunk mode.
A virtual node (VN) port is a port on an end host, such as a host bus adapter (HBA) port, not a port on an FC switch. It is this type of port to which VF ports are typically connected. Although a VN port might participate in an Ethernet VLAN that is using STP, in this scenario you have been asked to identify a switch port type for which an STP configuration might be a consideration.
A switched port analyzer (SPAN) port, which is also known as a mirroring port, is a type of port that is used to collect copies of packets transmitted over another port, over a given device, or over a network. In an FCoE configuration, a SPAN destination port can be either an FC interface or an Ethernet interface. SPAN source ports, on the other hand, can be FC interfaces, virtual FC (vFC) interfaces, a virtual SAN (vSAN), a VLAN, an Ethernet interface, a port channel interface, or a SAN port channel interface.

Reference:

Cisco: Fibre Channel over Ethernet Operations: FCoE and Spanning Tree Protocol Considerations

QUESTION 54
Which of the following is not a benefit of server virtualization?

A. reduces network bandwidth hot spots
B. can aid configuration standardization
C. reduces facility expenses
D. can grow and shrink based on resource need
E. reduces maintenance downtime

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Server virtualization causes, not reduces, network bandwidth hot spots. Therefore, reducing network bandwidth hot spots is not a benefit of server virtualization.
Server virtualization is the process of using a virtual machine (VM) hypervisor to create and maintain multiple server VMs on a single hardware server. Because all the VMs are using the same physical network connection, it is possible that required network bandwidth could surpass network capacity. It is therefore important that VM administrators understand the network demands of each VM installed on a given physical server and migrate or deploy new VMs on new hardware as necessary.
Server virtualization can reduce maintenance downtime. Because server VMs can be migrated from one hypervisor to another without powering down the VM, maintenance on a hypervisor's host need not hinder virtualized server availability.
Server virtualization can grow and shrink based on resource need. This feature is known as elasticity and is commonly used by cloud-based virtual servers. A virtual server that requires more hardware resources can grow to consume those resources and shrink when those resources are no longer required.
Server virtualization can reduce facility expenses. When servers are virtualized, it is not necessary to purchase new hardware for each new server that is deployed at the facility. Instead, new servers can be quickly instantiated on existing hardware, thereby eliminating the cost of purchasing additional hardware.
Server virtualization can aid configuration standardization. Because VMs can be cloned, administrators can create what is known as a golden image, which is a single VM that is equipped with a standardized, secure configuration. This golden image can then be cloned to create new VMs that are already equipped with that standard, secure configuration. Deploying and instantiating VMs that are already configured with a standard, secure configuration reduces administrative overhead and prevents accidental deployment of an insecure configuration.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 11: Server Virtualization Solutions, Server Virtualization Challenges, pp. 416-421

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 11: Server Virtualization Solutions, Server Virtualization Benefits, pp. 415-416

QUESTION 55

Examine the Cisco UCS Director Workflow Designer workflow in the following exhibit:

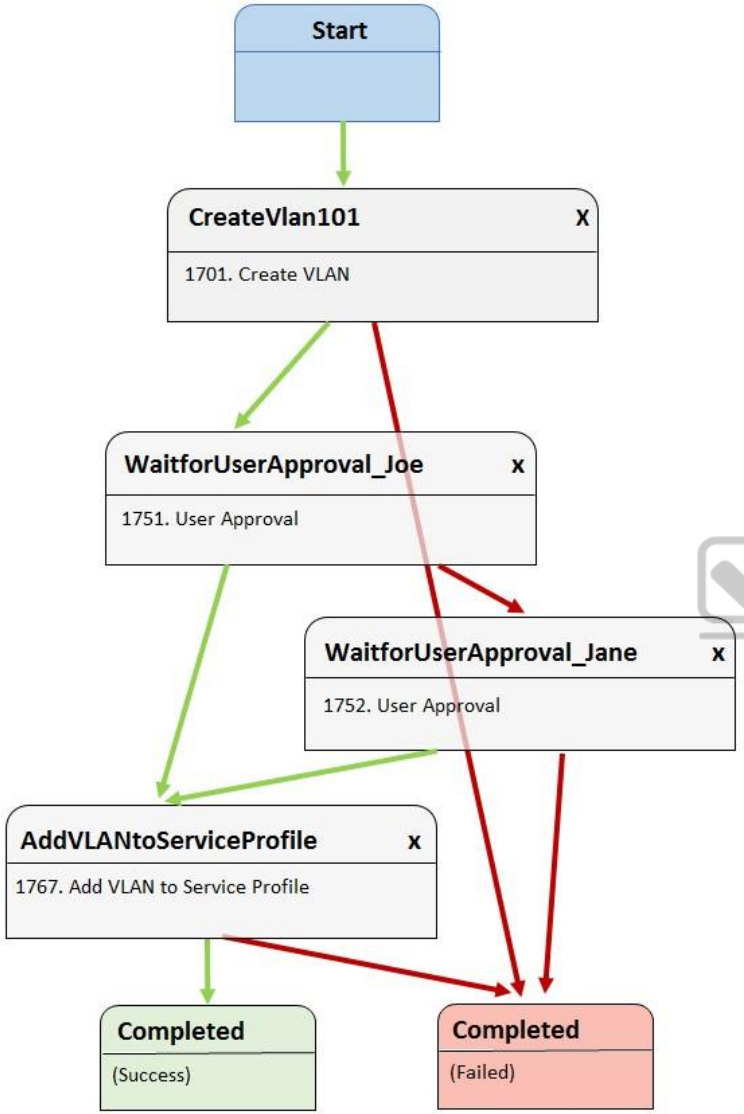

How many users must approve before VLAN 101 can be added to the service profile?

A. one
B. three
C. two
D. none

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

In this scenario, only one user must approve before virtual local area network (VLAN) 101 can be added to the service profile. The Cisco Unified Computing System (UCS) Director Workflow Designer workflow in this scenario creates VLAN 101 after it starts. Before the next task is executed, the workflow requires the approval of a user named Joe. If Joe approves, the workflow attempts to add VLAN 101 to the service profile. If Joe does not approve, Jane is given an opportunity to approve. If Jane approves, VLAN 101 can be added to the service profile.
Cisco UCS Director Workflow Designer is a graphical user interface (GUI) that enables users to create automated workflows in a drag-and-drop fashion. Each task in a workflow is equipped with an **On Success** button and an **On Failure** button. Each button provides a drop-down list of other tasks in the workflow. In this way, the user can select which tasks are executed next if a task succeeds and which tasks are executed next if a task fails. Green arrows in Workflow Designer represent the **On Success** path. Red arrows represent the **On Failure** path.
At least one user will need to approve in this scenario for VLAN 101 to be added to the service profile. The **On Failure** path of the **WaitforUserApproval_Joe** task is tied to the **WaitforUserApproval_Jane** task. The **On Failure** path of the **WaitforUserApproval_Jane** task is tied to the **Completed (Failure)** task. Therefore, there is no chance of successfully adding VLAN 101 to the service profile without user approval.
No more than one user will need to approve in this scenario for VLAN 101 to be added to the service profile. The **On Success** path of the **WaitforUserApproval_Joe** task is tied to the **AddVLANtoServiceProfile** task. Therefore, the workflow attempts to add VLAN 101 to the service profile if Joe approves. Similarly, the **On Success** path of the **WaitforUserApproval_Jane** task is tied to the **AddVLANtoServiceProfile** task. Therefore, the workflow attempts to add VLAN 101 to the service profile if Joe rejects and Jane approves.
Only two user approval tasks are present in this workflow: **WaitforUserApproval_Joe** and **WaitforUserApproval_Jane**. Therefore, it is not possible to require more than two users to approve or reject the addition of VLAN 101 to the service profile.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 17: Understanding and Troubleshooting UCSD Workflows, Creating Workflows, pp. 641-645

QUESTION 56
Which of the following examples best describes the PaaS service model?

A. A company licenses an office suite, including email service, that is delivered to the end user through a web browser.
B. A company obtains a subscription to use a service provider's infrastructure, programming tools, and programming languages to develop and serve cloud-based applications.
C. A company moves all company-wide policy documents to an Internet-based virtual file system hosted by a service provider.
D. A company hires a service provider to deliver cloud-based processing and storage that will house multiple virtual hosts configured in a variety of ways.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

A company that uses a service provider's infrastructure, programming tools, and programming languages to develop and serve cloud-based applications is an example of the Platform as a Service (PaaS) service model. The National Institute of Standards and Technology (NIST) defines three service models in its definition of cloud computing: Software as a Service (SaaS), Infrastructure as a Service (IaaS), and PaaS. Cloud computing offers several benefits over traditional physical infrastructure and software licensing, including a reduction in downtime and administrative overhead.
The SaaS service model enables its consumer to access applications running in the cloud infrastructure but does not enable the consumer to manage the cloud infrastructure or the configuration of the provided applications. A company that licenses a service provider's office suite and email service that is delivered to end users through a web browser is using SaaS. SaaS providers use an Internet-enabled licensing function, a streaming service, or a web application to provide end users with software that they might otherwise install and activate locally. Web-based email clients, such as Gmail and Outlook.com, are examples of SaaS. The PaaS service model provides its consumer with a bit more freedom than the SaaS model by enabling the consumer to install and possibly configure providersupported applications in the cloud infrastructure. A company that uses a service provider's infrastructure, programming tools, and programming languages to develop and serve cloud-based applications is using PaaS. PaaS enables a consumer to use the service provider's development tools or Application Programming Interface (API) to develop and deploy specific cloud-based applications or services. Another example of PaaS might be using a third party's MySQL database and Apache services to build a cloud-based customer relationship management (CRM) platform.
The IaaS service model provides the greatest degree of freedom by enabling its consumer to provision processing, memory, storage, and network resources within the cloud infrastructure. The IaaS service model also enables its consumer to install applications, including operating systems (OSs) and custom applications. However, with IaaS, the cloud infrastructure remains in control of the service provider. A company that hires a service provider to deliver cloud-based processing and storage that will house multiple physical or virtual hosts configured in a variety of ways is using IaaS. For example, a company that wanted to establish a web server farm by configuring multiple Linux Apache MySQL PHP (LAMP) servers could save hardware costs by virtualizing the farm and using a provider's cloud service to deliver the physical infrastructure and bandwidth for the virtual farm. Control over the OS, software, and server configuration would remain the responsibility of the organization, whereas the physical infrastructure and bandwidth would be the responsibility of the service provider.
A company that moves all company-wide policy documents to an Internet-based virtual file system hosted by a third party is using cloud storage. Cloud storage is a term used to describe the use of a service provider's virtual file system as a document or file repository. Cloud storage enables an organization to conserve storage space on a local network. However, cloud storage is also a security risk in that the organization might not have ultimate control over who can access the files.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 15: Cloud Computing, Cloud Computing Service Models, pp. 574-577

QUESTION 57
Which of the following best describes a port that operates as part of a FabricPath network?

A. a trunk port
B. an edge port
C. an access port
D. a core port

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, a core port is a port that operates as part of a FabricPath network. Cisco FabricPath uses Open Systems Interconnection (OSI) networking model Layer 3 routing combined with Layer 2 switching to construct a unified and scalable Layer 2 fabric. Although Cisco FabricPath defines two types of ports, only core ports are considered to be part of the FabricPath network. Core ports forward Ethernet frames encapsulated within a FabricPath header. In addition, core ports are always trunk ports that include an Institute of Electrical and Electronics Engineers (IEEE) 802.1Q virtual local area network (VLAN) tag. Only FabricPath VLANs are allowed on core ports.
An edge port is a Cisco FabricPath component port that does not operate as part of the FabricPath network. Instead, edge ports send only normal Ethernet frames as part of a classic Layer 2 switched network. An edge port can be configured as either an access port or an IEEE 802.1Q trunk port.
Although all core ports are trunk ports, not all trunk ports are core ports. Therefore, of the available choices, the term trunk port does not best describe a port that operates as part of a FabricPath network. Trunk ports enable switches to transmit and receive data on multiple VLANs over the same link.
An access port does not best describe a port that operates as part of a FabricPath network. Core ports cannot be access ports. An access port is a switch port that typically connects to an end device, such as a server or workstation. Access ports transmit and receive data on a single VLAN.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 1: Data Center Networking, Components of FabricPath, pp. 32-34

QUESTION 58
Which of the following statements about VE ports is true?

A. They connect to peripheral devices.
B. They enable multihop FCoE topologies.
C. They cannot be configured with multiple vSANs on a single port.

D. They are enabled in access mode by default.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Because virtual edge (VE) ports connect Fibre Channel over Ethernet (FCoE) inter-switch links (ISLs), VE ports enable multihop FCoE topologies. An FCoE hop is a connection from one domain ID to another domain ID. Therefore, a multihop FCoE switch is aware of multiple domains and can track and forward between them. In other words, VE ports connect to other VE ports.
VE ports do not connect to peripheral devices. Virtual F (VF) ports, on the other hand, are FCoE ports that typically connect to virtual N (VN) ports. VN ports are ports in peripheral devices, such as end hosts or disks.
VE ports are enabled in trunk mode by default. In addition, VE ports can be configured with multiple virtual storage area networks (vSANs) on a single port. In order to configure multiple vSANs on a single VE port, you must configure each corresponding FCoE virtual local area network (VLAN) on the Ethernet interface to which the VE port is bound.

Reference:

Cisco: Configuring Fibre Channel Interfaces: VE Ports
Cisco: Nexus 5500 to Nexus 7000 Multi-Hop FCoE Configuration Example: Introduction

QUESTION 59
Which of the following Cisco UCS Manager identity pools contains both types of FC1 identities?

A. the MAC pool
B. the IP pool
C. the WWNN pool
D. the WWxN pool
E. the WWPN pool

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, only the Cisco Unified Computing System (UCS) WWxN identity pool contains both Fibre Channel-Layer 2 (FC1) identities. This is because WWxN identity pools are a combination of World Wide Name (WWN) types: the World Wide Port Name (WWPN) and the World Wide Node Name (WWNN). Unlike other networking technologies, Fibre Channel (FC) does not make use of the Open Systems Interconnection (OSI) network model. Instead, FC uses the FC-Layers model, which is broken out in the following fashion:

- FC4: Protocol mapping layer
- FC3: Common services layer
- FC2: Network layer
- FC1: Data link layer
- FC0: Physical layer

FC0, FC1, and FC2 have similar names and functions to their OSI model equivalents, which are Layer 1, Layer 2, and Layer 3, respectively. FC3 is equivalent to the OSI model's Transport layer, or Layer 4. FC4, on the other hand, is similar in function to a combination of all three top layers of the OSI model, which are the Session layer (Layer 5), the Presentation layer (Layer 6), and the Application layer (Layer 7).

WWNNs are 64-bit globally unique identifiers that specify a given FC node. These identifiers are typically used to assign FC1 addresses in storage area network (SAN) routing. Similar to the WWNN identity pool, the WWPN identity pool contains globally unique 64-bit identifiers that are used to assign FC1 addresses. However, WWPNs represent a specific FC port, not an entire node.

Media Access Control (MAC) identity pools contain MAC addresses, which are OSI Layer 2 48-bit hexadecimal addresses that are typically burned into a network interface card (NIC). The first 24 bits of a MAC address represent the Organizationally Unique Identifier (OUI), which is a value that is assigned by the Institute of Electrical and Electronics Engineers (IEEE). The OUI identifies the NIC's manufacturer. The last 24 bits of a MAC address uniquely identify a specific NIC constructed by the manufacturer. This value is almost always an identifier that the manufacturer has never before used in combination with the OUI.

Internet Protocol (IP) identity pools contain IP addresses, which are 32-bit decimal addresses that are assigned to OSI Layer 3 interfaces. In a Cisco UCS domain, IP pools are typically used to assign one or more management IP addresses to each server's Cisco Integrated Management Controller (IMC).

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 9: Cisco Unified Computing System Pools, Policies, Templates, and Service Profiles, Cisco UCS Logical Resource Pools, pp. 340-341
IETF: RFC 4172: iFCP - A Protocol for Internet Fibre Channel Storage Networking: 3.3. Fibre Channel Layers and Link Services
Cisco: Overview of Cisco Unified Computing System: Pools

QUESTION 60

Which of the following statements about contracts in a Cisco ACI fabric is true?

A. Multicast traffic is not permitted among EPGs without a contract.
B. Members of an EPG require contracts in order to communicate with other members.
C. Contracts consist of subjects, filters, actions, and objects.
D. EPGs communicate with each other according to contract rules.

Correct Answer: D

Section: (none)
Explanation

Explanation/Reference:
Explanation:

Endpoint groups (EPGs) in a Cisco Application Centric Infrastructure (ACI) fabric communicate with each other according to contract rules. EPGs are logical groupings of endpoints that provide the same application or components of an application. For example, a collection of Hypertext Transfer Protocol Secure (HTTPS) servers could be logically grouped into an EPG labeled WEB. Contracts are policy objects that define how EPGs communicate. There are three types of contracts that can be applied in an ACI fabric:
Regular – applies filters to matching traffic and typically follows taboo contracts
Taboo – denies and logs matching traffic
Out-of-Band (OOB) - applies to OOB traffic from the management tenant
With the exception of some types of traffic—such as network configuration traffic, routing protocol traffic, and multicast traffic—EPGs require contracts in order to communicate with each other.
Members of an EPG do not require contracts in order to communicate with other members. Instead, members of an EPG communicate with each other by using their own network configurations, rules, and filters.
Contracts consist of subjects, filters, actions, and optionally labels, not objects. Subjects are groups of filters that are specific to a given application. Filters classify traffic by matching Open Systems Interconnection (OSI) network model Layer 2 or Layer 4 characteristics. Actions are the action that is performed on traffic that matches the filters. Labels can be created to group EPGs or subjects. These groupings add granularity to the enforcement of a policy.
Multicast traffic is permitted among EPGs without a contract. In addition, some Dynamic Host Configuration Protocol version 4 (DHCPv4) traffic is permitted between EPGs without a contract. Other traffic types that are permitted between EPGs by default are Open Shortest Path First (OSPF), Enhanced Interior Gateway
Routing Protocol (EIGRP), Protocol Independent Multicast (PIM), Internet Group Management Protocol (IGMP), and Internet Control Message Protocol version 6 (ICMPv6) neighbor discovery.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 13: ACI Logical Model and Policy Framework, Contracts, pp. 506-508
Cisco: Working with Contracts: Contracts

QUESTION 61
Which of the following is a benefit of CNAs in a Cisco Unified Fabric?

A. consistent policies
B. elimination of STP
C. reduced cabling
D. segregation of LAN and SAN

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, reduced cabling is a benefit of converged network adapters (CNAs) in a Cisco Unified Fabric. Cisco Unified Fabric is a combination of architecture and high performance concepts that is intended to simplify data center networks.
CNAs are network adapters that combine network interface cards (NICs) and host bus adapters (HBAs), enabling one adapter to support both Ethernet and Fibre Channel (FC). A server that contains separate FC and local area network (LAN) ports can require significantly more cabling than a server that is configured with a CNA.
Segregation of LAN and storage area network (SAN) is not a benefit of CNAs in a Cisco Unified Fabric. On the contrary, CNAs are a feature of Cisco Unified Fabric that help converge a data center's LAN and SAN over a single transport in order to simplify management, provisioning, and operation.
Elimination of Spanning Tree Protocol (STP) is not a benefit of CNAs in a Cisco Unified Fabric. Virtual Port Channels (vPCs) reduce Cisco Unified Fabric's reliance on STP by replacing EtherChannel. However, vPCs still require STP to mitigate switching loops if they occur. Cisco FabricPath can replace STP with the Layer 3 routing protocol Intermediate System-to-Intermediate System (IS-IS) in order to scale a Layer 2 network beyond normal limits.
Consistent policies are not a benefit of CNAs in a Cisco Unified Fabric. Consistent policies across Cisco Unified Fabric is a benefit of using Cisco NX-OS and policy templates that can be easily deployed across the fabric and in virtual environments. These templates reduce the likelihood of human error during configuration and ensure that security and performance are applied the same way across the network.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 3: Unified Fabric Overview, Convergence of Network and Storage, pp. 82-88

QUESTION 62
You want to configure IP pools and MAC pools in the Cisco UCS Manager GUI.
Which of the following Navigation pane tables should you click?

A. **Servers**
B. **SAN**
C. **VM**
D. **Admin**
E. **LAN**
F. **Equipment**

Correct Answer: E
Section: (none)

Explanation

Explanation/Reference:
Explanation:

You should click the **LAN** tab in the Cisco Unified Computing System (UCS) Manager graphical user interface (GUI) if you want to configure Internet Protocol (IP) pools and Media Access Control (MAC) pools. Cisco UCS Manager GUI is a Java application. The main area of the GUI is divided into a Navigation pane and a work area. The selections you click in the Navigation pane determine the information and configuration fields that are displayed in the work area.
The Navigation pane of the Cisco UCS Manager GUI contains six tabs in a row across the top of the pane. To configure or view a given element of Cisco UCS Manager, you should first click the Navigation pane tab appropriate for that element. The Navigation pane contains all of the following tabs: ▪ The **Equipment** tab
▪ The **Servers** tab
▪ The **LAN** tab
▪ The **SAN** tab
▪ The **VM** tab
▪ The **Admin** tab
The **Equipment** tab can be used to display an inventory of the UCS domain. This information includes color-coded fault indicators, such as a red, yellow, or orange rectangle. If a device has a fault, one of these indicators will appear around the name of the device on the **Equipment** tab. The **Equipment** tab contains four nodes: **Equipment**, **Chassis**, **Rack-mounts**, and **Fabric Interconnects**. Selected nodes contain information specific to the devices indicated by the node name. For example, the **Fabric Interconnects** node contains information about expansion modules, fans, and power supply units (PSUs) connected to the domain's fabric interconnects.
The **Servers** tab can be used to modify server-specific configurations, such as policies, profiles, and universally unique identifier (UUID) pools. The **Servers** tab contains six nodes: **Servers**, **Service Profiles**, **Service Profile Templates**, **Policies**, **Pools**, and **Schedules**. Selected nodes contain information specific to the server configurations indicated by the node name. For example, the **Policies** node allows the configuration of policies related to server adapters, server firmware, and other components.
The **LAN** tab can be used to configure local area network (LAN) components, such as Quality of Service (QoS) classes, virtual LANs (VLANs), and flow control policies. The **LAN** tab contains seven nodes: **LAN Cloud**, **Appliances**, **Internal LAN**, **Policies**, **Pools**, **Traffic Monitoring Sessions**, and **Netflow Monitoring**. Selected nodes contain information specific to the LAN component indicated by the node name. For example, the **Pools** node allows the configuration of both IP address pools and MAC address pools that have been defined for a LAN.
The **SAN** tab can be used to configure storage area network (SAN) components, such as virtual SANs (vSANs), and World Wide Name (WWN) pools. The **SAN** tab contains six nodes: **SAN**, **SAN Cloud**, **Storage Cloud, Policies**, **Pools**, and **Traffic Monitoring Sessions**. Selected nodes contain information specific to the SAN component indicated by the node name. For example, the **SAN** node allows the configuration of SAN uplinks, Fibre Channel (FC) address assignments, and vSANs.
The **VM** tab can be used to configure virtual machine-fabric extender (VM-FEX) for UCS domain servers that are equipped with virtual interface cards (VICs). The **VM** tab contains seven nodes: **All**, Clusters, **Fabric Network Sets**, **Port Profiles**, **VM Networks**, **Microsoft**, and **VMware**. Selected nodes contain information specific to the VM component indicated by the node name. For example, the **VMware** node can be used to configure Cisco UCS Manager connections to VMware vCenter.
The **Admin** tab can be used to configure system-wide settings that must be configured by an administrator or viewed by a security administrator. The **Admin** tab contains 10 nodes:

- **All**
- **Faults, Events and Audit Log**
- **User Management**
- **Key Management**
- **Communication Management**
- **Stats Management**
- **Time Zone Management**
- **Capability Catalog**
- **Management Extension**
- **License Management**

Selected nodes contain information specific to the administrative component indicated by the node name. For example, the **User Management** node allows the configuration of authentication methods and user roles as well as remote access methods.

Reference:

Cisco: Overview of Cisco UCS Manager GUI: Navigation Pane

QUESTION 63

You are configuring a Cisco Nexus 5000 Series switch for the first time.
Which of the following is true about the default login?

A. You can configure a short, trivial password at first boot and change it later.
B. You will be required to configure an admin password before configuration.
C. The default user is the network admin, and the password is **password**.
D. You are automatically logged in as a low-level local user.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

When configuring a Cisco Nexus 5000 Series switch for the first time, you will be required to configure an admin password before configuration. This step in the configuration process is required and cannot be skipped by using the Ctrl-C keyboard combination. After you have successfully configured an admin password, you can enter setup mode by entering **yes** at the prompt.
The default user is the network admin; however, the default password is not **password**. When configuring a Cisco Nexus 5000 Series switch for the first time, the admin account is not configured with a password. By default, the network admin account is named **admin**. This account cannot be changed or deleted. However, the account password can and must be set at first boot.

You cannot configure a short, trivial password at first boot and change it later. If you attempt to configure a short, trivial password at boot, the Cisco Nexus 5000 Series configuration script will reject the password.
You will not be automatically logged in as a low-level local user. When a Cisco Nexus 5000 Series switch boots for the first time, only one default account, named admin, exists on the switch. After you have configured a strong password for the admin account, you can create additional user accounts on the switch. However, the switch must be initially configured by using the admin account.

Reference:

Cisco: Initial Switch Configuration: Default Login

QUESTION 64
Which of the following network devices do not connect to a leaf switch?

A. a spine switch
B. a router
C. another leaf switch
D. an APIC controller

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, only another leaf switch does not connect to a leaf switch. In a spine-leaf architecture, the leaf layer of switches provides connectivity to and scalability for all other devices in the data center network. However, leaf switches do not connect to other leaf switches. Instead, a leaf switch communicates with another leaf switch by using a spine switch. The following exhibit displays a typical spine-leaf architecture wherein the top row of devices represents the spine layer of switches and the bottom row of devices represents the leaf layer of switches:

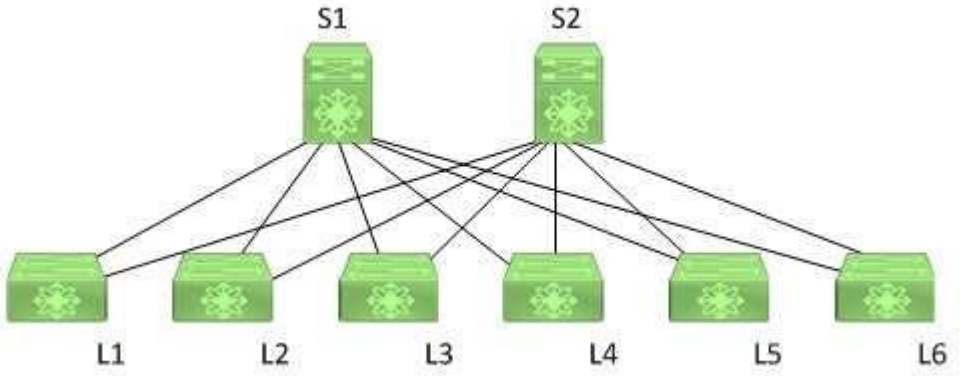

In the exhibit above, the leaf switches are each directly connected to both spine switches. Switches S1 and S2 comprise the spine layer of the topology. Switches L1, L2, L3, L4, L5, and L6 comprise the leaf layer of the topology. In order for L1 to send a packet to L6, the packet must traverse either S1 or S2. The spine-leaf architecture differs from the traditional three-tier network architecture, which consists of a core layer, an aggregation layer, and an access layer.
Leaf switches connect to spine switches. In a spine-leaf architecture, spine switches are used to provide bandwidth and redundancy for leaf switches. Therefore, spine switches do not connect to devices other than leaf switches. As the name implies, spine switches are the backbone of the architecture.
Leaf switches connect to Cisco Application Policy Infrastructure Controllers (APICs). The leaf switch to which a Cisco APIC is directly connected is the first device in a spine-leaf architecture that will be discovered by and registered with the APIC. When a Cisco APIC begins the switch discovery process, it first detects only the leaf switch to which it is connected. After that leaf switch is registered, the APIC discovers each of the spine switches to which the leaf switch is connected. Spine switches do not automatically register with the APIC. When a spine switch is registered with the APIC, the APIC will discover all the leaf switches that are connected to that spine switch.
Leaf switches connect to routers. Routers are typically used to connect to the Internet or to a wide area network (WAN). Leaf switches in a spine-leaf architecture directly connect to routers.
Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 12: ACI Architecture, Spine-Leaf Data Center Design, pp. 447-449

QUESTION 65
You connect Ethernet links to four ports on slot 1 of a new Cisco Nexus 5548UP switch and receive the following error:

```
ERROR: Ethernet range starts from first port of the module
ERROR: FC range should end on last port of the module
```

Which of the following is most likely true?

A. Slot 1 does not support unified ports.
B. No more Ethernet ports are available on the module.

C. The Ethernet links were connected in an incorrect port order.
D. FC links are already connected to slot 1.
E. Slot 1 does not support Ethernet links.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Most likely, the Ethernet links were connected in an incorrect port order if you receive the error message in this scenario when you connect Ethernet links to four ports on slot 1 of a new Cisco Nexus 5548UP switch. Unified ports are ports that can be either Ethernet or Fibre Channel (FC) ports. Cisco Nexus switches that support unified ports require that Ethernet links be connected from the beginning of the port range forward and that FC links be connected from the end of the port range backward. Therefore, in this scenario, the Ethernet links would need to be connected to port 1/1, port 1/2, port 1/3, and port 1/4.
As a further example, if you were to connect a single Ethernet link and a single FC link to slot 1 of a Cisco Nexus 5548UP switch that has no other links connected, you must connect the Ethernet link to port 1/1 and the FC link to port 1/32. Slot 1 of a Cisco Nexus 5548UP switch contains 32 unified ports. If you were to then connect an additional Ethernet link to slot 1, that link must be connected to port 1/2. If you were to connect an additional FC link to slot 1, that link must be connected to port 1/31. Link additions should continue in that way until the Ethernet range ends where the FC range begins and no more ports are available in the slot.
There is nothing in this scenario to indicate that FC links are already connected to slot 1 of the Cisco Nexus 5548UP switch. However, even if FC links were already connected to slot 1 of the switch, the error would still indicate that the Ethernet links were connected in an incorrect port order.
It is not likely that there are no more Ethernet ports available on the switch. In this scenario, you have already connected four Ethernet links to the switch when the error message occurs. There would have to have been at least five Ethernet links available on the switch in this scenario in order for the switch to have detected the incorrect port order.
Slot 1 on the Cisco Nexus 5548UP switch supports Ethernet links because slot 1 on the Cisco Nexus 5548UP switch contains only unified ports. Other Nexus switches might contain different port configurations. For example, on the Cisco Nexus 5596T, only the last 16 ports are unified ports.

Reference:
Cisco: Configuring Layer 2 Interfaces: Information About Unified Ports

QUESTION 66
You connect a FEX to a Cisco Nexus switch. After configuration is complete, you issue the **show fex 100** command on the switch and receive the following partial output:

```
Switch#show fex 100
<output omitted>
Pinning-mode: static Max-links: 1
Fabric port for control traffic: Po100
Fabric interface state:
```

```
Po100 - Interface Up. State: Active
Eth1/1 - Interface Up. State: Active
Eth1/2 - Interface Up. State: Active
Eth2/1 - Interface Up. State: Active
Eth2/2 - Interface Up. State: Active
```

Which of the following is true?

A. The **fex associate Eth2/1** command has been issued on the Po100 interface.
B. The **fex associate Eth2/2** command has been issued on the Po100 interface.
C. The **fex associate Eth1/1** command has been issued on the Po100 interface.
D. The **fex associate 100** command has been issued on the Po100 interface.
E. The **fex associate Eth1/2** command has been issued on the Po100 interface.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

The **fex associate 100** command has been issued on the Po100 port-channel interface in this scenario. There are two methods for connecting a Cisco fabric extender (FEX) to a Cisco Nexus switch: static pinning and port channel. In this scenario, the output of the **show fex 100** command on the Nexus switch indicates that four Nexus switch physical interfaces and one port channel have been associated with the FEX: Ethernet 1/1, Ethernet 1/2, Ethernet 2/1, Ethernet 2/2, and P0100. In addition, the Po100 port channel is acting as the FEX fabric port.
To associate a port channel with a FEX, you should first create the port channel, configure the port channel to support a FEX, and then associate the port channel with a specific FEX. For example, the following commands configure Po100 on a Cisco Nexus switch to support a FEX and associate with FEX 100:
interface port-channel 100
switchport mode fex-
fabric fex associate 100
The **switchport mode fex-fabric** command and the **fex associate 100** command should also be issued on each physical port that is a member of Po100's channel group 100.
None of the other commands in this scenario could be issued on the Cisco Nexus switch, because they all contain invalid syntax. The syntax of the **fex associate** command is **fex associate** FEX-number, where *FEX-number* is the FEX chassis ID that is assigned to the FEX. Chassis IDs are assigned by the administrator and are valid in the range from 100 through 199. Chassis ID values less than 100 are typically ports on the parent switch. Ports on a FEX device are identified by chassis number, slot number, and port number. For example, a FEX interface might be identified as Ethernet 100/1/1, which indicates the interface in FEX 100 slot 1, port 1.

Reference:

Cisco: Configuring the Fabric Extender: Associating a Fabric Extender to a Port Channel

QUESTION 67
Which of the following Cisco ACI logical constructs is equivalent to a private IP name space or IP network?

A. common
B. EPG
C. context
D. tenant

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, a context is a Cisco Application Centric Infrastructure (ACI) fabric logical construct that is equivalent to a single Internet Protocol (IP) network or IP name space. In this way, a context can be considered equivalent to a single virtual routing and forwarding (VRF) instance. An endpoint policy typically ensures that all endpoints within a given context exhibit similar behavior.
Tenants are containers that can be used to represent physical tenants, organizations, domains, or specific groupings of information. A tenant in a Cisco ACI fabric can therefore contain multiple contexts. Typically, tenants are configured to ensure that different policy types are isolated from each other, similar to user groups or roles in a role-based access control (RBAC) environment.
Common is the name for a special tenant in a Cisco ACI fabric. The common tenant typically contains policies that can be shared with other tenants. Contexts that are placed in the common tenant can likewise be shared among tenants. Contexts that are placed within a private tenant, on the other hand, are not shared with other tenants.
Endpoint groups (EPGs) are logical groupings of endpoints that provide the same application or components of an application. For example, a collection of Hypertext Transfer Protocol Secure (HTTPS) servers could be logically grouped into an EPG labeled WEB. EPGs are typically collected within application profiles. EPGs can communicate with other EPGs by using contracts.

Reference:
Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 13: ACI Logical Model and Policy Framework, ACI Logical Constructs, pp. 488-492

QUESTION 68
Which of the following statements about vPC technology limitations is not true?

A. Each VDC is a separate switch.
B. There is only one vPC domain ID per switch.
C. There are only two vPC domain IDs per switch.

D. vPC peer links are always 10 Gbps.
E. vPC forms a Layer 3 port channel.

Correct Answer: E
Section: (none)
Explanation

Explanation/Reference:
Explanation:

A virtual port channel (vPC) forms an Open System Interconnection (OSI) networking model Layer 2 port channel, not a Layer 3 port channel. The vPC feature is not capable of supporting Layer 3 port channels. Therefore, any routing that is configured from the vPC peers to other parts of the network should be performed on separate Layer 3 ports.
Each virtual device context (VDC) is a separate switch in a vPC domain. A VDC logically virtualizes a switch. A VDC is a single virtual instance of physical switch hardware. A vPC logically combines ports from multiple switches into a single port-channel bundle. Conventional port channels, which are typically used to create high-bandwidth trunk links between two switches, require that all members of the bundle exist on the same switch. vPCs enable virtual domains that are comprised of multiple physical switches to connect as a single entity to a fabric extender, server, or other device.
A vPC domain is comprised of two switches per domain. In addition, a vPC domain cannot be comprised of more than two switches. Each switch in the vPC domain must be configured with the same vPC domain ID. To enable vPC configuration on a Cisco Nexus 7000 Series switch, you should issue the **feature vpc** command on both switches. To assign the vPC domain ID, you should issue the **vpc domain** *domain-id* command, where *domain-id* is an integer in the range from 1 through 1000, in global configuration mode. For example, issuing the **vpc domain 101** command on a Cisco Nexus 7000 Series switch configures the switch with a vPC domain ID of 101.
Only one vPC domain can be configured per switch. If you were to issue more than one **vpc domain** *domain-id* command on a Cisco Nexus 7000 Series switch, the vPC domain ID of the switch would become whatever value was issued last. After you issue the **vpc domain** *domain-id* command, the switch is placed into vPC domain configuration mode. In vPC domain configuration mode, you should configure a peer keepalive link.
Peer keepalive links monitor the remote device to ensure that it is operational. You can configure a peer keepalive link in any virtual routing and forwarding (VRF) instance on the switch. Each switch must use its own Internet Protocol (IP) address as the peer keepalive link source IP address and the remote switch's IP address as the peer keepalive link destination IP address. The following commands configure a peer keepalive link between SwitchA and SwitchB in vPC domain 101:

```
SwitchA(config)#vpc domain 101
SwitchA(config-vpc-domain)#peer-keepalive destination 192.168.1.2 source 192.168.1.1 vrf default
SwitchB(config)#vpc domain 101
SwitchB(config-vpc-domain)#peer-keepalive destination 192.168.1.1 source 192.168.1.2 vrf default
```

A vPC peer link should always be a 10-gigabit-per-second (Gbps) Ethernet port, not 1 Gbps. Peer links are configured as a port channel between the two members of the vPC domain. You should configure vPC peer links after you have successfully configured a peer keepalive link. Cisco recommends connecting two 10-Gbps Ethernet ports from two different input/output (I/O) modules. To configure a peer link, you should issue the **vpc peer-link** command in interface configuration mode. For example, the following commands configure a peer link on port-channel 1

```
SwitchA(config)#interface port-channel 1
```

```
SwitchA(config-if)#switchport mode trunk SwitchA(config-
if)#vpc peer-link
SwitchB(config)#interface port-channel 1
SwitchB(config-if)#switchport mode trunk
SwitchB(config-if)#vpc peer-link
```

It is important to issue the correct **channel-group** commands on a port channel's member ports prior to configuring the port channel. For example, if you are creating Port-channel 1 by using the Ethernet 2/1 and Ethernet 2/2 interfaces, you could issue the following commands on each switch to correctly configure those interfaces as members of the port channel:

```
SwitchA(config)#interface range ethernet 2/1-2
Switch(config-if-range)#switchport
SwitchA(config-if-range)#channel-group 1 mode active
SwitchB(config)#interface range ethernet 2/1-2
SwitchB(config-if-range)#switchport
SwitchB(config-if-range)#channel-group 1 mode active
```

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 1: Data Center Networking, vPC Limitations, p. 27

QUESTION 69
Which of the following best describes the control plane of a Nexus switch?

A. It is where SNMP operates.
B. It is where routing calculations are made.
C. It is also known as the forwarding plane.
D. It is where traffic forwarding occurs.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, the control plane of a Nexus switch is best described as where routing calculations are made. A Nexus switch consists of three operational planes: the control plane, the management plane, and the data plane, which is also known as the forwarding plane. Routing protocols such as Enhanced Interior Gateway Routing Protocol (EIGRP), Open Shortest Path First (OSPF), and Border Gateway Protocol (BGP) all operate in the control plane of a Nexus switch. The control plane is responsible for gathering and calculating the information required to make the decisions that the data plane needs for

forwarding. Routing protocols operate in the control plane because they enable the collection and transfer of routing information between neighbors. This information is used to construct routing tables that the data plane can then use for forwarding.
The data plane, not the control plane, is where traffic forwarding occurs. Cut-through switching allows a switch to begin forwarding a frame before the frame has been received in its entirety. Store-and-forward switching receives an entire frame and stores it in memory before forwarding the frame to its destination.
The management plane, not the control plane, is where Simple Network Management Protocol (SNMP) operates. The management plane is responsible for monitoring and configuration of the control plane. Therefore, network administrators typically interact directly with protocols running in the management plane.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 2: Management and Monitoring of Cisco Nexus Devices, Operational Planes of a Nexus Switch, pp. 50-59

QUESTION 70
Which of the following protocols are used to bring up an ACI fabric in a cascading manner? (Choose two.)

A. LLDP
B. CDP
C. STP
D. DHCP
E. ARP

Correct Answer: AD
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Link-Layer Discovery Protocol (LLDP) and Dynamic Host Configuration Protocol (DHCP) are protocols that are used to bring up a Cisco Application Centric Infrastructure (ACI) fabric in a cascading manner. LLDP is a standard protocol that detects neighboring devices of any type. An ACI fabric uses LLDP along with DHCP to discover switch nodes and to assign Internet Protocol (IP) addresses to virtual extensible local area network (VXLAN) tunnel endpoints (VTEPs). LLDP is also used by Cisco Application Policy Infrastructure Controllers (APICs) to detect virtual switches, although it is also possible to use Cisco Discovery Protocol (CDP) for that purpose.
Each APIC server in the ACI communicates with ACI nodes and other APIC servers during the APIC cluster discovery process by using private IP addresses. Private IP addresses are not routable over the Internet. When an ACI fabric is booted, an internal private IP addressing scheme is used to enable the APIC to communicate with other nodes and controllers.
The ACI fabric does not use CDP, Address Resolution Protocol (ARP), or Spanning Tree Protocol (STP) to bring up a Cisco API fabric in a cascading manner. If configured, CDP can be used instead of LLDP to detect virtual switches. ARP is a protocol that is, used to map Media Access Control (MAC) addresses to IP addresses on a LAN. STP is a protocol that is used in switched networks to prevent switching loops.

Reference:

Cisco: Cisco Application Centric Infrastructure Fundamentals: Startup Discovery and Configuration

QUESTION 71
Which of the following is also known as a context within an ACI tenant?

A. an EPG
B. a bridge domain
C. a VRF instance
D. an application profile

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, a virtual routing and forwarding (VRF) instance is also known as a context within a Cisco Application Centric Infrastructure (ACI) tenant. VRF instances are also sometimes known as private networks. A VRF instance is an Open Systems Interconnection (OSI) networking model Layer 3 forwarding domain within a given ACI tenant. Multiple bridge domains can be connected to a given VRF instance within a tenant.
A bridge domain typically defines a Layer 2 address space and flood domain within a Cisco ACI fabric. An ACI bridge domain is similar to a traditional networking virtual local area network (VLAN) in that it is a Layer 2 broadcast domain. Bridge domains define the Media Access Control (MAC) address space. Typically, a bridge domain is associated with a single subnet, although multiple subnets can be associated within a given bridge domain. Bridge domains are typically connected to VRF instances within a given ACI tenant.
An application profile is used to configure policies and relationships between endpoint groups (EPGs). Typically, an application profile is a container for one or more logically related EPGs. For example, EPGs that provide similar services or functions might be associated with the same application profile.
An EPG is a logical ACI construct that contains multiple related endpoints. Endpoints are physical or virtual devices that are connected to a network. For example, a web server is an endpoint. EPGs contain endpoints that have similar policy requirements, although not necessarily similar functions. EPGs enable group management of the policies for the endpoints they contain.

Reference:

Cisco: Cisco Application Centric Infrastructure Fundamentals: VRFs

QUESTION 72

Which of the following logically virtualizes an OSI networking model Layer 3 address domain?

A. a VRF instance

B. vPC
C. a VIC
D. a VDC

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

A virtual routing and forwarding (VRF) instance logically virtualizes an Open Systems Interconnection (OSI) networking model Layer 3 address domain. VRFs enable a router to maintain multiple, simultaneous routing tables. Therefore, VRFs can be configured on a single router to serve multiple Layer 3 domains instead of implementing multiple hardware routers.
A virtual device context (VDC) does not logically virtualize an OSI networking model Layer 3 address domain. A VDC logically virtualizes a Cisco Nexus 7000 Series switch. A VDC is a single virtual instance of physical switch hardware. By default, the control plane of the Cisco Nexus 7000 Series switch is configured to run a single VDC. It is possible to configure multiple VDCs on the same hardware. A single VDC can contain multiple virtual local area networks (VLANs) and VRF instances.
A virtual port channel (vPC) does not logically virtualize an OSI networking model Layer 3 address domain. A vPC enables ports from multiple switches to be combined into a single port channel bundle. Conventional port channels, which are typically used to create high-bandwidth trunk links between two switches, require that all members of the bundle exist on the same switch. vPCs enable virtual domains that are comprised of multiple physical switches to connect as a single entity to a fabric extender, server, or other device.
A virtual interface card (VIC) does not logically virtualize an OSI networking model Layer 3 address domain. A VIC is a Cisco device that can be used to create multiple logical network interface cards (NICs) and host bus adapters (HBAs). VICs such as the Cisco M81KR send Fibre Channel over Ethernet (FCoE) traffic and normal Ethernet traffic over the same physical medium.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 6: Virtualizing Cisco Network Devices, Describing Layer 3 Virtualization Within VDCs, pp. 204-206

QUESTION 73
You connect Port 13 of a Cisco UCS Fabric Interconnect to a SAN. All the Fabric Interconnect ports are unified ports.
Which of the following options should you select for Port 13 in Cisco UCS Manager?

A. **Configure as Appliance Port**
B. **Configure as Server Port**
C. **Configure as FCoE Storage Port**
D. **Configure as Uplink Port**
E. **Configure as FCoE Uplink Port**

Correct Answer: E
Section: (none)
Explanation

Explanation/Reference:
Explanation:

You should select the Cisco Unified Computing System (UCS) Manager **Configure as FCoE Uplink Port** option for Port 13 on the Cisco UCS Fabric Interconnect in this scenario because Port 13 is connected to a storage area network (SAN). Cisco UCS Fabric Interconnects enable management of two different types of network fabrics in a single UCS domain, such as a Fibre Channel (FC) SAN and an Ethernet local area network (LAN).
Unified ports in a UCS Fabric Interconnect can operate in one of two primary modes: FC or Ethernet. Each mode can be assigned a different port role, such as server port, uplink port, or appliance port for Ethernet. The FC port mode supports two roles: FC over Ethernet (FCoE) uplink port and FCoE storage port. Uplink ports are used to connect the Fabric Interconnect to the next layer of the network, such as to a SAN by using the FCoE uplink port role or to a LAN by using the uplink port role.
You should not choose the **Configure as Uplink Port** option in this scenario. The **Configure as Uplink Port** option would configure Port 13 in the uplink role in Ethernet mode. This port role is used for connecting a UCS Fabric Interconnect to an Ethernet LAN.
You should not choose the **Configure as Server Port** option in this scenario. The **Configure as Server Port** option is used to connect the Fabric Interconnect to network adapters on server hosts.
You should not choose the **Configure as FCoE Storage Port** option in this scenario. The **Configure as FCoE Storage Port** option is used to connect the Fabric Interconnect to a Direct-Attached Storage (DAS) device.
You should not choose the **Configure as Appliance Port** option in this scenario. The **Configure as Appliance Port** option is used to connect the Fabric Interconnect to a storage appliance.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 8: Cisco UCS Manager, Basic Port Roles in the Cisco UCS Fabric Interconnects, pp. 318-319

QUESTION 74
Which of the following Cisco FabricPath features improves on data center scalability?

A. use of iSCSI instead of FCoE
B. deployment of consistent policies

C. elimination of reliance on STP
D. convergence of network and storage

Correct Answer: C
Section: (none)
Explanation
Explanation/Reference:
Explanation:

Of the available choices, the elimination of reliance on Spanning Tree Protocol (STP) to ensure a loop-free switching environment is a Cisco FabricPath feature that improves data center scalability and growth. Cisco Unified Fabric uses virtual Port Channel (vPC) in place of technologies such as EtherChannel, which was developed to enable redundant high-speed connectivity between switches in an STP topology. However, STP is still present to ensure that switching loops can be mitigated if they occur. Cisco FabricPath, on the other hand, is a Cisco Unified Fabric technology that completely replaces STP with the Intermediate System-toIntermediate System (IS-IS) routing protocol. The combination of IS-IS with the Open Systems Interconnection (OSI) networking model Layer 2 fabric's simplicity and fabric extenders enhances the scalability of Cisco Unified Fabric beyond the practical limits of a normal Layer 2 topology.
The deployment of consistent network policies is a Cisco Unified Fabric feature that improves on data center security, not a Cisco FabricPath feature that improves on data center scalability. Cisco Unified Fabric allows the use of templates and a common switch operating system (OS) to ensure the deployment of network policies consistently across the fabric and its virtualized environments. The use of templates reduces the likelihood of human error when deploying network policies.
In addition, Cisco Unified Fabric contains virtualization-aware security products.
The use of Internet Small Computer Systems Interface (iSCSI) instead of Fibre Channel over Ethernet (FCoE) in a Cisco Unified Fabric enables the encapsulation of Fibre Channel (FC) in Transmission Control Protocol/Internet Protocol (TCP/IP) packets; it is not a Cisco FabricPath feature that improves on data center scalability. The use of iSCSI in a Cisco Unified Fabric can be considered an alternative to the use of FCoE in a fabric that does not have strict storage connectivity requirements. Unlike iSCSI, FCoE encapsulates FC in Ethernet frames.
Convergence of network and storage is a Cisco Unified Fabric feature that simplifies operation and reduces management endpoints, not a Cisco FabricPath feature that improves on data center scalability. A typical Cisco Unified Fabric architecture is used to merge storage area network (SAN) features with a local area network (LAN). The resulting converged network and storage is delivered over an Ethernet fabric.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 3: Unified Fabric Overview, Scalability and Growth, pp. 88-93

QUESTION 75
Which of the following port types allows only FabricPath VLAN traffic?

A. a trunk port
B. a core port
C. an access port
D. an edge port

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, only a Cisco FabricPath core port allows only FabricPath virtual local area network (VLAN) traffic. Cisco FabricPath uses Open Systems Interconnection (OSI) networking model Layer 3 routing combined with Layer 2 switching to construct a unified and scalable Layer 2 fabric. Although Cisco FabricPath defines two types of ports, only core ports are considered to be part of the FabricPath network. Core ports forward Ethernet frames encapsulated within a FabricPath header. In addition, core ports are always trunk ports that include an Institute of Electrical and Electronics Engineers (IEEE) 802.1Q VLAN tag. Only FabricPath VLANs are allowed on core ports.
Classic Ethernet VLAN traffic is allowed on an edge port. An edge port is a Cisco FabricPath component port that does not operate as part of the FabricPath network. Instead, edge ports send only normal Ethernet frames as part of a classic Layer 2 switched network. An edge port can be configured as either an access port or an IEEE 802.1Q trunk port.
Although all core ports are trunk ports, not all trunk ports are core ports. Therefore, a trunk port could carry either FabricPath VLAN traffic or classic Ethernet VLAN traffic depending on the circumstances and configuration. Trunk ports enable switches to transmit and receive data on multiple VLANs over the same link.
An access port carries either FabricPath VLAN traffic or classic Ethernet VLAN traffic, depending on the circumstances and configuration. An access port is a switch port that typically connects to an end device, such as a server or workstation. Access ports transmit and receive data on a single VLAN. Core ports cannot be access ports.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 1: Data Center Networking, Components of FabricPath, pp. 32-34

QUESTION 76
Examine the Cisco UCS Director Workflow Designer workflow in the following exhibit:

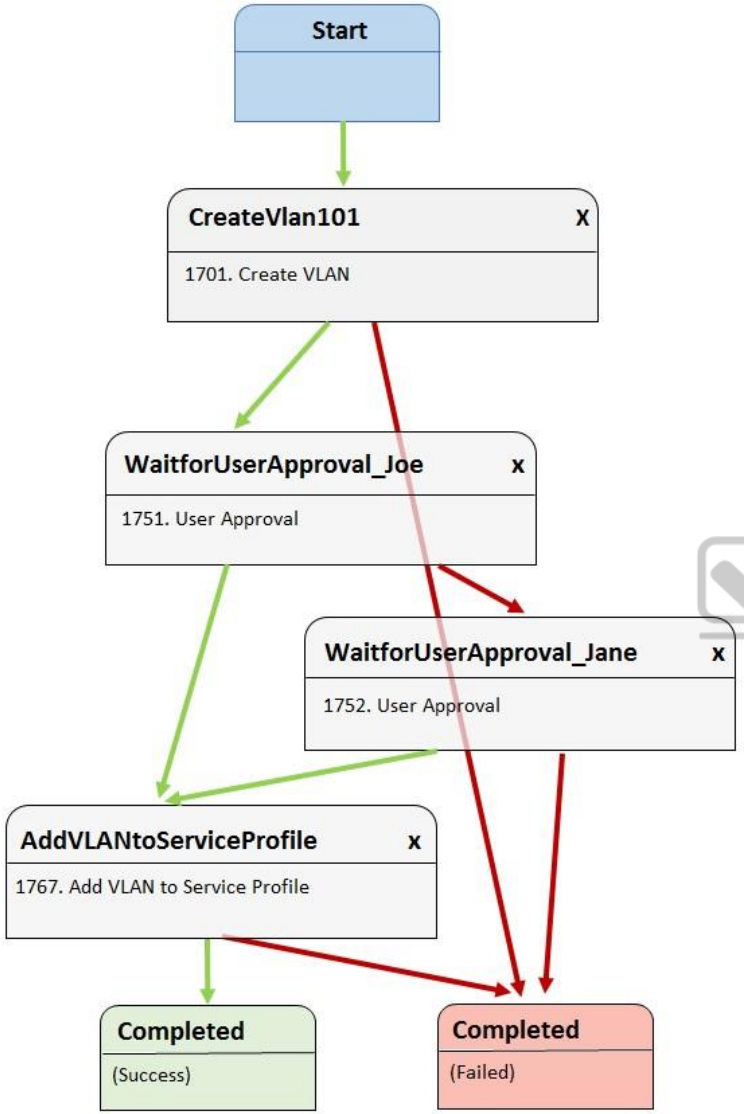

Jane initiates the workflow, but Joe rejects it.

Which of the following best describes what will most likely occur next?

A. VLAN 101 will be sent back to Joe for approval if Jane approves.
B. VLAN 101 will not be added to the service profile, and the workflow will succeed.
C. VLAN 101 will not be added to the service profile, and the workflow will fail.
D. VLAN 101 will be added to the service profile if Jane approves.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Most likely, virtual local area network (VLAN) 101 will be added to the service profile if Jane approves. The Cisco Unified Computing System (UCS) Director Workflow Designer workflow in this scenario creates VLAN 101 after it starts. Before the next task is executed, the workflow requires the approval of a user named Joe. In this scenario, Joe has rejected the approval step. If Joe's approval is not obtained, then the workflow prompts Jane for approval to continue. If Jane rejects the approval step, the workflow ends in failure. If Jane approves, VLAN 101 will be assigned to the service profile.
Cisco UCS Director Workflow Designer is a graphical user interface (GUI) that enables users to create automated workflows in a drag-and-drop fashion. Each task in a workflow is equipped with an **On Success** button and an **On Failure** button. Each button provides a drop-down list of other tasks in the workflow. In this way, the user can select which tasks are executed next if a task succeeds and which tasks are executed next if a task fails. Green arrows in Workflow Designer represent the **On Success** path. Red arrows represent the **On Failure** path.
VLAN 101 will not be sent back to Joe for approval if Jane approves. In this scenario, the **WaitforUserApproval_Jane** task's **On Success** button is connected to the **AddVLANtoServiceProfile** task. Therefore, Jane can approve adding VLAN 101 to the service profile even if Joe does not.
There is nothing in this scenario that would lead to workflow failure or workflow success based on Joe's rejection of the task, because Jane has an opportunity to approve the addition of the VLAN to the Service Profile if Joe rejects it. However, if Joe were to approve the addition of VLAN 101 to the service profile, the task would be attempted and potentially completed without Jane's approval or rejection because the **WaitforUserApproval_Joe** task's On Success button is tied directly to the **AddVLANtoServiceProfile** task.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 17: Understanding and Troubleshooting UCSD Workflows, Creating Workflows, pp. 641-645

QUESTION 77
Which of the following are benefits of server virtualization? (Choose four.)

A. reduces maintenance downtime
B. reduces facility expenses

C. can aid configuration standardization
D. reduces network bandwidth hot spots
E. can grow and shrink based on resource need

Correct Answer: ABCE
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Server virtualization can reduce maintenance downtime. Server virtualization is the process of using a virtual machine (VM) hypervisor to create and maintain multiple server VMs on a single hardware server. Because server VMs can be migrated from one hypervisor to another without powering down the VM, maintenance on a hypervisor's host need not hinder virtualized server availability.
Server virtualization can grow and shrink based on resource need. This feature is known as elasticity and is commonly used by cloud-based virtual servers. A virtual server that requires more hardware resources can grow to consume those resources and shrink when those resources are no longer required.
Server virtualization can reduce facility expenses. When servers are virtualized, it is not necessary to purchase new hardware for each new server that is deployed at the facility. Instead, new servers can be quickly instantiated on existing hardware, thereby eliminating the cost of purchasing additional hardware.
Server virtualization can aid configuration standardization. Because VMs can be cloned, administrators can create what is known as a golden image, which is a single VM that is equipped with a standardized, secure configuration. This golden image can then be cloned to create new VMs that are already equipped with that standard, secure configuration. Deploying and instantiating VMs that are already configured with a standard, secure configuration reduces administrative overhead and prevents accidental deployment of an insecure configuration.

Server virtualization causes, not reduces, network bandwidth hot spots. Therefore, reducing network bandwidth hot spots is not a benefit of server virtualization. Because all the VMs are using the same physical network connection, it is possible that required network bandwidth could surpass network capacity. It is therefore important that VM administrators understand the network demands of each VM installed on a given physical server and migrate or deploy new VMs on new hardware as necessary.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 11: Server Virtualization Solutions, Server Virtualization Benefits, pp. 415-416
Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 11: Server Virtualization Solutions, Server Virtualization Challenges, pp. 416-421

QUESTION 78
You are installing a Cisco Nexus 1000v VSM in VMware vSphere by using the recommended OVA file method of installation.
Which of the following steps are performed automatically? (Choose four.)

A. install VSM plug-in
B. create the VSM

C. add hosts
D. configure VSM networking
E. configure SVS connection
F. perform initial VSM setup

Correct Answer: ABDF
Section: (none)
Explanation

Explanation/Reference:
Explanation:

The creation of the virtual supervisor module (VSM), the configuration of VSM networking, the initial VSM setup, and the installation of the VSM plug-in are performed automatically when you install a Cisco Nexus 1000v VSM in VMware vSphere by using the open virtualization appliance (OVA) file method of installation. You are required to manually configure the server virtualization switch (SVS) connection and add hosts. The OVA file method of installation is the Cisco-preferred method of installing a Nexus 1000v VSM on VMware vSphere.
There are three methods of installing the Nexus 1000v VSM: ▪
By using an International Standards Organization (ISO) image
▪ By using an open virtualization format (OVF) folder and installation wizard ▪
By using an OVA file and installation wizard
No matter which installation method is chosen, the VSM installation process consists of the following six steps: ▪
Creation of the VSM virtual machine (VM) in Cisco vCenter
▪ Configuration of VSM networking
▪ Initial VSM setup in the VSM console
▪ Installation of the VSM plug-in in vCenter
▪ Configuration of the SVS connection in the VSM console
▪ Addition of hosts to the virtual distributed switch in vCenter
The steps in this process that you are required to perform manually depend on the method of installation you choose.
The OVF folder method of installing the Nexus 1000v VSM is similar to the OVA file method. However, the OVF folder method performs only the first two steps of the installation process. This means that you are required to manually configure the final four steps. The primary difference between an OVA file and an OVF folder is that the OVA file is a single compressed archive.
The ISO image method of installing the Nexus 1000v VSM requires that you manually perform each of the six steps in the installation process. An ISO image file contains a virtual filesystem that can be mounted by an operating system (OS) similar to an optical disc or a Universal Serial Bus (USB) flash drive.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 4: Cisco Nexus 1000V and Virtual Switching, Cisco Nexus 1000V VSM Installation Methods, pp. 144-145

QUESTION 79
Which of the following Cisco VIC adapter communication methods is most likely to rely on the VMware ESXi hypervisor to switch traffic?

A. software and Nexus 1000V
B. pass-through switching
C. bare-metal OS driver
D. store-and-forward switching

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, the Cisco virtual interface card (VIC) adapter method that is most likely to rely on the VMware ESXi hypervisor to switch traffic is software and a Cisco Nexus 1000v virtual switch. Cisco VICs integrate with virtualized environments to enable the creation of virtual network interface cards (vNICs) in virtual machines (VMs). When using software and a Cisco Nexus 1000v virtual switch to switch traffic, switching is controlled by the hypervisor. Cisco describes the software-based method of handling traffic by using a Cisco Nexus 1000v virtual switch as Virtual Network Link (VN-Link).
Pass-through switching, which Cisco describes as hardware-based VN-Link, does not rely on the VMware ESXi hypervisor to switch traffic. Pass-through switching is a faster and more efficient means for Cisco VIC adapters to handle traffic between VMs.
Pass-through switching uses application-specific integrated circuit (ASIC) hardware switching, which reduces overhead because the switching occurs in the fabric instead of relying on software. Pass-through switching also enables administrators to apply network policies between VMs in a fashion similar to how traffic policies are applied between traditional physical network devices.
A bare-metal operating system (0S) driver method of communication is not likely to rely on the VMware ESXi hypervisor to switch traffic. Cisco VIC adapters forward traffic similar to other Cisco Unified Computing System (UCS) adapters when installed in a server that is configured with a single bare-metal OS without virtualization. It is therefore possible to use a Cisco VIC adapter with OS drivers to create static vNICs on a server in a nonvirtualized environment.
Store-and-forward switching is not a method of communication that is used by Cisco VIC adapters. However, switches can be configured to use store-and-forward switching. A switch that uses store-and-forward switching receives the entire frame before forwarding the frame. By receiving the entire frame, the switch can verify that no cyclic redundancy check (CRC) errors are present in the frame; this helps prevent the forwarding of frames with errors. Cisco VIC adapters are hardware interface components that support virtualized network environments.

Reference:

Cisco: Overview of VN-Link in Cisco UCS

QUESTION 80
Which of the following are not capable of operating as a standalone physical switch? (Choose two.)

A. Nexus 5000 Series
B. Nexus 2000 Series
C. Nexus 7000 Series
D. Nexus 9000 Series
E. Nexus 1000v Series

Correct Answer: BE
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, neither the Cisco Nexus 1000v Series nor the Cisco Nexus 2000 Series are capable of operating as a standalone physical switch. The Cisco Nexus 1000v is a virtual switch that is capable of connecting to upstream physical switches in order to provide connectivity for a virtual machine (VM) network environment. Although the Cisco Nexus 1000v operates similar to a standard switch, it exists only as software in a virtual environment and is therefore not a physical switch.
The Cisco Nexus 2000 Series of switches are fabric extenders (FEXs) and cannot operate as standalone switches. FEX technologies depend on parent switches, such as a Cisco Nexus 5500 Series switch or a Cisco Nexus 7000 Series switch, to provide forwarding tables and control plane functionality. FEX technologies are intended to extend the network to edge devices. Typically, FEX devices in the Cisco Nexus 2000 Series are managed by first connecting to the parent device using either Telnet or Secure
Shell (SSH) and then configuring the FEX.
Cisco Nexus 5000 Series switches operate as standalone physical switches. Cisco Nexus 5000 Series switches are data center access layer switches that can support 10-gigabit-per-second (Gbps) or 40-Gbps Ethernet, depending on the model. Native Fibre Channel (FC) and FC over Ethernet (FCoE) are also supported by Cisco Nexus 5000 Series switches.
Cisco Nexus 7000 Series switches operate as standalone physical switches. Cisco Nexus 7000 Series switches are typically used as an end-to-end data center solution, which means that the series is capable of supporting all three layers of the data center architecture: core layer, aggregation layer, and access layer. In addition, the Cisco Nexus 7000 Series supports virtual device contexts (VDCs). The Cisco Nexus 7000 Series can support up to 100-Gbps Ethernet.
Cisco Nexus 9000 Series switches operate as standalone physical switches. Cisco Nexus 9000 Series switches can operate either as traditional NX-OS switches or in an Application Centric Infrastructure (ACI) mode. Unlike Cisco Nexus 7000 Series, Cisco Nexus 9000 Series switches do not support VDCs or storage protocols.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 4: Cisco Nexus 1000V and Virtual Switching, Cisco Nexus 1000V System Overview, p. 134
Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 6: Virtualizing Cisco Network Devices, Nexus 2000 Series Fabric Extender Connectivity, pp. 211-212

QUESTION 81

Which of the following are most likely to operate in the management plane of a Nexus switch?

A. store-and-forward switching
B. cut-through switching
C. EIGRP
D. SNMP
E. OSPF
F. BGP

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, Simple Network Management Protocol (SNMP) is an Internet Protocol (IP) network management protocol that is most likely to operate in the management plane of a Nexus switch. A Nexus switch consists of three operational planes: the data plane, which is also known as the forwarding plane, the control plane, and the management plane. The management plane is responsible for monitoring and configuration of the control plane. Therefore, network administrators typically interact directly with protocols running in the management plane.
Cut-through switching and store-and-forward switching are most likely to operate in the data plane of a Nexus switch. Of the three, the data plane is where traffic forwarding occurs. Cut-through switching allows a switch to begin forwarding a frame before the frame has been received in its entirety. Store-and-forward switching receives an entire frame and stores it in memory before forwarding the frame to its destination.
Enhanced Interior Gateway Routing Protocol (EIGRP), Open Shortest Path First (OSPF), and Border Gateway Protocol (BGP) all operate in the control plane of a Nexus switch. The control plane is responsible for gathering and calculating the information required to make the decisions that the data plane needs for forwarding. Routing protocols operate in the control plane because they enable the collection and transfer of routing information between neighbors. This information is used to construct routing tables that the data plane can then use for forwarding.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 2: Management and Monitoring of Cisco Nexus Devices, Management Plane, pp. 5859

QUESTION 82
You issue the following command on a Cisco Nexus 7000 switch:
vrf context management
Which of the following is most likely to occur?

A. The switch will be placed into VRF configuration mode for an existing VRF instance.
B. The switch will return an error because the management VRF already exists.

C. A new VDC named **management** will be created.
D. A new VRF instance named **management** will be created.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Most likely, the switch will be placed into virtual routing and forwarding (VRF) configuration mode for the existing management VRF. There are two VRF instances configured on a Nexus 7000 by default: the management VRF and the default VRF. The management VRF is used only for management, includes only the mgmt 0 interface, and uses only static routing.
The **vrf context** *name* command can be used to create a new VRF or to enter VRF configuration mode for an existing VRF. Because the Nexus 7000 is already configured with a management VRF, issuing the command in this scenario places the device into VRF configuration mode for that VRF. Similarly, issuing the **vrf context default** command would place the switch into VRF configuration mode for the existing default VRF.
VRFs are used to logically separate Open Systems Interconnection (OSI) networking model Layer 3 networks. Therefore, it is possible to have overlapping Internet Protocol version 4 (IPv4) or Internet Protocol version 6 (IPv6) addresses in environments that contain multiple tenants. However, an interface that has been assigned to a given VRF cannot be simultaneously assigned to another VRF. The address space, routing process, and forwarding table that are used within a VRF are local to that VRF.
The default VRF, on the other hand, includes all Layer 3 interfaces until you assign those interfaces to another VRF. Similarly, the default VRF runs any routing protocols that are configured unless those routing protocols are assigned to another VRF. All **show** and **exec** commands that are issued in the default VRF apply to the default routing context. Unless an administrator configures other VRFs on a Nexus 7000, any forwarding configurations that are made by the administrator will operate in the default VRF.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 6: Virtualizing Cisco Network Devices, Describing Layer 3 Virtualization Within VDCs, pp. 204-206
Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 13: ACI Logical Model and Policy Framework, ACI Logical Constructs, pp. 488-492

QUESTION 83
You want to perform maintenance on an ESXi server and need to migrate its VMware VMs to another ESXi server. However, you do not need to migrate the VM datastore. You do not want to power off the VMs during the migration process.

Which of the following solutions should you choose?

A. vSphere vMotion
B. copying or cloning
C. cold migration

D. Storage vMotion

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, you should choose VMware vMotion if you want to migrate a VMware virtual machine (VM) from one VMware ESXi server to another ESXi server without powering off the VM. In this scenario, you do not want to migrate the VMware datastore. VMware's vSphere vMotion does not migrate datastores. VMware's ESXi server is a bare-metal server virtualization technology, which means that ESXi is installed directly on the hardware it is virtualizing instead of running on top of another operating system (OS). This layer of hardware abstraction enables tools like vMotion to migrate ESXi VMs from one host to another without powering off the VM, enabling the VM's users to continue working without interruption.
You do not need to choose Storage vMotion to perform the migration in this scenario, because Storage vMotion allows migration of both the VM and its datastore. The datastore is the repository of VM-related files, such as logs and virtual disks. When migrating a VM by using Storage vMotion, both the virtualized environment and the datastore can be moved to a new host without powering down or suspending the VM.
You should not choose cold migration, copying, or cloning in this scenario. Cold migration is the process of powering down a VM and moving the VM or the VM and its datastore to a new location. While a cold migration is in progress, no users can perform tasks inside the VM. Both copying and cloning create new instances of a given VM. Therefore, neither action is a form of migrating a VM to another host. Typically, a VM must be powered off or suspended in order to successfully copy or clone it.

Reference:

VMware: VMware Docs: Migrating Virtual Machines
VMware: Virtualization Overview

QUESTION 84
You are examining the following command-line output on a Nexus 7000 Series switch:

```
vdc_idvdc_name state    mac
---------     -------- ------    ---------------
2             acctg     active    7F:52:92:26:29:CF
```

Which of the following have you most likely issued?

A. the **show vdc detail** command in the default VDC
B. the **show vdc** command in the default VDC
C. the **show vdc membership** command in a nondefault VDC
D. the **show vdc detail** command in a nondefault VDC

E. the **show vdc** command in a nondefault VDC

Correct Answer: E
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Most likely, you have issued the show vdc command in a nondefault virtual device context (VDC) if you see the command-line output in this scenario on a Nexus 7000 Series switch. The **show vdc** command displays information about VDCs configured on the physical switch. However, the output of the **show vdc** command depends on the VDC in which the command has been issued. When issued in the default VDC, the command displays output for all VDCs configured on the device, as shown in the following output:

```
NEX7S#show vdc
vdc_id        vdc_name state     mac
---------     -------- ------    ---------------
1             salesactive      A8:CC:D8:D8:65:F4
2             acctgactive      7F:52:92:26:29:CF
3             prod active      7B:05:CA:ED:41:E1
```

When issued in a nondefault VDC, the command displays output for only the current VDC, as shown in the following output:

```
NEX7S#show vdc

vdc_id        vdc_name state     mac
---------     -------- ------    ---------------
2             acctg      active    7F:52:92:26:29:CF
```

To issue read and write commands in the default VDC, a user must be assigned the network-admin user role. The network-operator role has read-only access to the default VDC. Users that have been assigned the vdc-admin role or the vdc-operator role can review output from the **show vdc** command in the nondefault VDC in which they are operating. However, they do not have rights to read or write information in other VDCs.
It is not likely that you have issued the **show vdc detail** command in this scenario. Similar to the **show vdc** command, the **show vdc detail** command displays information about the VDCs configured on the physical device. However, the **show vdc detail** command provides extra information about the VDCs that is not visible in the summarized **show vdc** command output. For example, the output of the **show vdc detail** command contains information about the VDC ha policy, boot order, create time, and restart count in addition to the VDC ID, name, state, and Media Access Control (MAC) address. Also similar to the **show vdc** command, the **show vdc detail** command displays output for all VDCs when the command is issued in the default VDC. When issued in a nondefault VDC, the command displays output only for the VDC in which it was issued.
It is not likely that you have issued the **show vdc membership** command in this scenario. The **show vdc membership** command displays the interfaces that have been allocated to VDCs on the physical device, as shown in the following output:

```
NEX7S#show vdc membership vdc_id:    2
vdc_name: acctg interfaces:
        Ethernet2/1
```

Based on the output above, you can surmise that the command was issued in a nondefault VDC. The output displays information for only the VDC named acctg. If the command had been issued in the default VDC, the output would have displayed information about the interfaces assigned to every VDC that is configured on the physical device.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 6: Virtualizing Cisco Network Devices, Verifying VDCs on the Cisco Nexus 7000 Series Switches, pp. 201-204

QUESTION 85
Which of the following describes a group of Layer 3 networks that can be shared with groups of private networks in a Cisco ACI fabric?

A. EPG
B. common
C. tenant
D. context

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, common best describes a group of Layer 3 networks that can be shared with groups of private networks in a Cisco Application Centric Infrastructure (ACI) fabric. Common is the name for a special tenant in a Cisco ACI fabric. The common tenant typically contains policies that can be shared with other tenants.
Tenants are containers that can be used to represent physical tenants, organizations, domains, or specific groupings of information. A tenant in a Cisco ACI fabric can contain multiple Layer 3 networks. Typically, tenants are configured to ensure that different policy types are isolated from each other, similar to user groups or roles in a role-based access control (RBAC) environment.
A context is a Cisco ACI fabric logical construct that is equivalent to a single Internet Protocol (IP) network or IP name space. In this way, a context can be considered equivalent to a single virtual routing and forwarding (VRF) instance. An endpoint policy typically ensures that all endpoints within a given context exhibit similar behavior. Contexts that are placed in the common tenant can likewise be shared among tenants. Contexts that are placed within a private tenant, on the other hand, are not shared with other tenants.
Endpoint groups (EPGs) are logical groupings of endpoints that provide the same application or components of an application. For example, a collection of Hypertext Transfer Protocol Secure (HTTPS) servers could be logically grouped into an EPG labeled WEB. EPGs are typically collected within application profiles. EPGs can communicate with other EPGs by using contracts.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 13: ACI Logical Model and Policy Framework, ACI Logical Constructs, pp. 488-492

QUESTION 86
Which of the following is not a type of contract that can be applied in an ACI fabric?

A. regular
B. in-band
C. OOB
D. taboo

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Of the available choices, in-band is not a type of contract that can be applied in a Cisco Application Centric Infrastructure (ACI) fabric. Contracts are policy objects that define how endpoint groups (EPGs) communicate. There are three types of contracts that can be applied in an ACI fabric: Regular — applies filters to matching traffic and typically follows taboo contracts
Taboo — denies and logs matching traffic
Out-of-Band (OOB) - applies to OOB traffic from the management tenant
EPGs are logical groupings of endpoints that provide the same application or components of an application. For example, a collection of Hypertext Transfer Protocol Secure (HTTPS) servers could be logically grouped into an EPG labeled WEB. With the exception of some types of traffic—such as network configuration traffic, routing protocol traffic, and multicast traffic—EPGs require contracts in order to communicate with each other.
Contracts consist of subjects, filters, actions, and optionally, labels. Subjects are groups of filters that are specific to a given application. Filters classify traffic by matching Open Systems Interconnection (OST) network model Layer 2 or Layer 4 characteristics. Actions are the action that is performed on traffic that matches the filters. Labels can be created to group EPGs or subjects. These groupings add granularity to the enforcement of a policy.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 13: ACI Logical Model and Policy Framework, Contracts, pp. 506-508
Cisco: Working with Contracts: Contracts

QUESTION 87
You issue the following commands on a Nexus 7000 Series switch that is already configured to authenticate users by using TACACS+:

switchto vdc myvdc
configure terminal aaa

user default-role no
aaa user default-role
exit copy running-config start-
config

Which of the following will occur when a remote user attempts to log in to the VDC named MyVDC by using TACACS+?

A. The user will be assigned the vdc-operator role.
B. The user will be assigned the network-operator role.
C. The user will be assigned the network-admin role.
D. The user will be assigned the vdc-admin role.
E. The user will not be assigned a role and will be denied login.

Correct Answer: E
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The user will not be assigned a role and will be denied login when the remote user attempts to log in to the virtual device context (VDC) named MyVDC by using Terminal Access Controller Access-Control System Plus (TACACS+). In this scenario, TACACS+ is already configured on the Cisco Nexus 7000 Series switch. In addition, the **aaa user default-role** command has been issued but is immediately followed by the **no aaa user default-role** command in the configuration. Remote users who attempt to log in to the VDC named MyVDC will be denied access because no user role is assigned to those users.
Cisco Nexus switches use role-based access control (RBAC) to assign management privileges to a given user. By default, a Nexus 7000 switch is configured with the following user roles:
▪ network-admin — has read and write access to all VDCs on the switch ▪ network-operator — has read-only access to all the VDCs on the switch ▪ vdc-admin — has read and write access to a specific VDC on the switch ▪ vdc-operator — has read-only access to a specific VDC on the switch
The user will not be assigned the vdc-operator role, because the **no aaa user default-role** command has been issued. In this scenario, the **aaa user default-role** command has been issued in the VDC named MyVDC, which is a nondefault VDC on the switch. The **aaa user default-role** command configures the Authentication, Authorization, and Accounting (AAA) feature on the switch to automatically assign remote users the default user role at login. The default remote user role for nondefault VDCs on a Cisco Nexus switch is the vdc-operator role. However, this configuration will not apply in this scenario because of the **no aaa user default-role** command.
The user will not be assigned the vdc-admin user role. The vdc-admin user role allows read and write access to a specific VDC on the switch. If remote users were automatically assigned the vdc-admin role when logging in to the VDC named MyVDC, those users would have administrative access to the VDC, which is a security risk.

The user will not be assigned the network-admin role. In addition, the user will not be assigned the network-operator role. These roles are applied to users who have access to all VDCs that are configured on the switch, not a specific nondefault VDC. If the **aaa user default-role** command had been issued in the default VDC in this scenario, remote users who log in to the default VDC would be assigned a network-operator user role.

Reference:

Cisco: Configuring AAA: Enabling the Default User Role for AAA Authentication

QUESTION 88
The application network profile is an expression of which Cisco ACI model implementation stage?

A. the logical model
B. the concrete model
C. the hardware model
D. the resolved model

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Of the available choices, the application network profile is an expression of the logical model stage of the Cisco Application Centric Infrastructure (ACI) model. The application network profile, which is often simply referred to as the application profile, is an object that acts as the container for all other objects in the model; it tiers endpoint groups (EPGs) and represents how they relate to each other. The logical model is typically configured by the user in the Cisco Application Policy Infrastructure Controller (APIC).
The application network profile is not an expression of the resolved model stage of the Cisco ACI model. The resolved model is a stage that is derived by the APIC in abstract from the logical model stage. In other words, the logical model is converted to the resolved model by the APIC. This stage is a lower-level representation of the ACI model than the logical model stage in that it represents how configuration components are sent to the infrastructure when a policy is executed. The application network profile is not an expression of the concrete model stage of the Cisco ACI model. The concrete model is a stage that is derived from the resolved model in that it represents the ACI model after configuration has been delivered to each endpoint in the fabric. The concrete model is typically created when switches in the ACI model convert the resolved model. However, there are some objects in the logical model that do not require conversion to the resolved model to also be represented in the concrete model. These objects include physical port properties that are already local to a given leaf or port. The application network profile is not an expression of the hardware model stage of the Cisco ACI model. The Cisco ACI model consists of only three implementation stages. The hardware model is not one of those stages. Instead, concrete objects are consumed by application-specific integrated circuits (ASICs) on hardware.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 13: ACI Logical Model and Policy Framework, Different Models, pp. 487-488
Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 13: ACI Logical Model and Policy Framework, Application Profiles, p. 506

QUESTION 89
You issue the **port-channel load-balance scr-dst-port module 4** command on a Cisco Nexus 7000 Series switch running NX-OS 5.0(5).
How is distribution loaded for port channels on slot 4?

A. based on the source IP address and destination port
B. based on the source and destination MAC address
C. based on the source and destination port
D. based on the source MAC address and destination port only
E. based on the defaults because the command contains invalid syntax

Correct Answer: E
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Distribution for port channels on slot 4 in this scenario will be based on the defaults because the command contains invalid syntax. In this scenario, the Cisco Nexus 7000 Series switch is running NX-OS 5.0(5). Prior to NX-OS 5.1(3), the **port-channel load-balance** command required an **ethernet** keyword. Therefore, the valid syntax for NX-OS versions older than 5.1(3) is **port-channel load-balance ethernet** *method* [**module** *slot*]. In NX-OS 5.1(3), Cisco removed the ethernet keyword.
After NX-S 5.1(3), the basic syntax of the **port-channel load-balance** command is **port-channel load-balance** *method* [**module** *slot*], where *method* is one of the following 10 keywords:
▪ **dst-ip** ▪ **dst-mac** ▪ **dst-port** ▪ **src-dst-ip** ▪ **src-dst-mac** ▪ **src-dst-port**
▪ **src-ip** ▪ **src-mac** ▪ **src-port** ▪ **vlan-only**
The **dst-ip** keyword, **dst-mac** keyword, and **dst-port** keyword load port channel distribution based on the destination Internet Protocol (IP) address, Media Access Control (MAC) address, and port number, respectively. Similarly, the **src-ip** keyword, **src-mac** keyword, and **src-port** keyword load port channel distribution based on the source IP address, source MAC address, and port number, respectively. The **src-dst-ip** keyword, **src-dst-mac** keyword, and **src-dst-port** keyword load port channel distribution based on the source and destination IP addresses, MAC addresses, and port numbers, respectively. The **vlan-only** keyword loads distribution on only the virtual local area network (VLAN) modules.

The optional **module** keyword accepts a slot number value. If you configure the **port-channel load-balance** command with the **module** keyword, the configuration applies only to the specified slot. Otherwise, the configuration applies to the entire device. By default on a Cisco Nexus switch, a port channel load balances Layer 2 packets based on the source and destination MAC addresses. Layer 3 packets, on the other hand, are load balanced based on the source and destination IP addresses. You must be operating in the default virtual device context (VDC) on the switch in order to issue the **port-channel load-balance** command.

Reference:

Cisco: Cisco Nexus 7000 Series NX-OS Interfaces Command Reference: port-channel load-balance

QUESTION 90
You administer the Cisco UCS domain in the following exhibit:

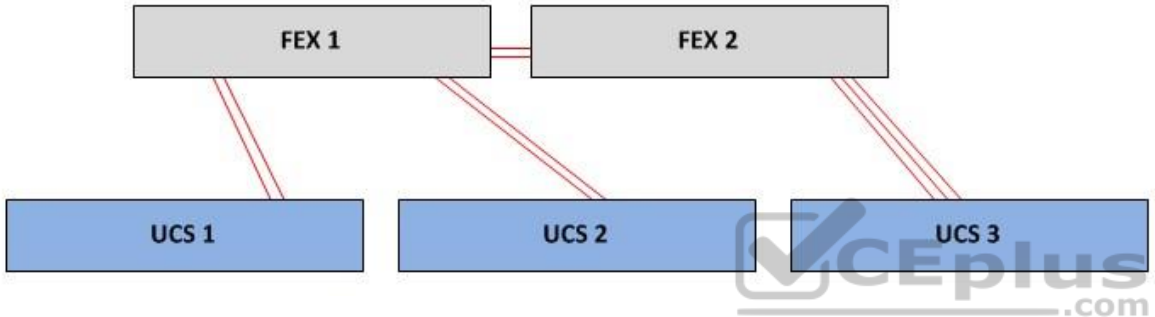

Which of the following policies could you implement to ensure that all three UCS chassis are discovered by Cisco UCS Manager? (Choose two.)
A. 4-Link Chassis Discovery Policy
B. 8-Link Chassis Discovery Policy
C. 3-Link Chassis Discovery Policy
D. 1-Link Chassis Discovery Policy
E. 2 Link Chassis Discovery Policy

Correct Answer: DE
Section: (none)
Explanation

Explanation/Reference:
Explanation:

You could implement either a 1-Link Chassis Discovery Policy or a 2-Link Chassis Discovery Policy to ensure that all three Cisco Unified Computing System (UCS) chassis are discovered by Cisco UCS Manager in this scenario. The link chassis discovery policy that is configured on Cisco UCS Manager specifies the minimum number of links that a UCS chassis must have to the fabric interconnect in order to be automatically discovered and added to the UCS domain. If a UCS

chassis has fewer than the minimum number of links specified by the policy, the chassis will be neither discovered nor added to the UCS domain by Cisco UCS Manager. The two Cisco fabric interconnects in this scenario, FEX 1 and FEX 2, are connected together. UCS 1 has a two-link connection to FEX 1. UCS 2 has a two-link connection to FEX 1. UCS 3 has a three-link connection to FEX 2.
A 1-Link Chassis Discovery Policy in this scenario would cause Cisco UCS Manager to discover and add UCS 1, UCS 2, and UCS 3 to the UCS domain because each of those devices has at least one link to the fabric interconnect. However, each of the devices
will be discovered and added as a chassis with only one link to the fabric interconnect. In order to have Cisco UCS Manager recognize and use the other links between the UCS devices and the fabric interconnect, you would need to reacknowledge the chassis in Cisco UCS Manager after the initial discovery has been completed.
A 2-Link Chassis Discovery Policy in this scenario would cause Cisco UCS Manager to discover and add UCS 1, UCS 2, and UCS 3 to the UCS domain because each of those devices has at least two links to the fabric interconnect. However, UCS 3, which has three links to the fabric interconnect, will be discovered and added as a chassis with only two links. In order to have Cisco UCS Manager recognize and use the other links between UCS 3 and the fabric interconnect, you would need to reacknowledge the chassis in Cisco UCS Manager after the initial discovery has been completed.
A 3-Link Chassis Discovery Policy would not cause Cisco UCS Manager to discover the UCS devices in this scenario, because Cisco UCS Manager does not support a 3-Link Chassis Discovery Policy. There are five link chassis discovery policies that are supported by Cisco UCS Manager: 1-Link Chassis Discovery Policy, 2-Link Chassis Discovery Policy, 4-Link Chassis Discovery Policy, 8-Link Chassis Discovery Policy, and Platform-Max Discovery Policy. The Platform-Max Discovery Policy is not a choice in this scenario.
A 4-Link Chassis Discovery Policy would not cause Cisco UCS Manager to discover the UCS devices in this scenario. For a 4-Link Chassis Discovery Policy to discover any of the UCS devices in this scenario, at least one of those devices would need to have a minimum of four links to the fabric interconnect. For example, if UCS 3 had four or more links to the fabric interconnect in this scenario and Cisco UCS Manager was configured with a 4-Link Chassis Discovery Policy, only UCS 3 would be discovered and added to the UCS domain.
An 8-Link Chassis Discovery Policy would not cause Cisco UCS Manager to discover the UCS devices in this scenario. For an 8-Link Chassis Discovery Policy to discover any of the UCS devices in this scenario, at least one of those devices would need to have a minimum of eight links to the fabric interconnect. None of the UCS devices in this scenario are connected with more than three links to the FEX devices.
Reference:

Cisco: Configuring System-Related Policies: Chassis Discovery Policy

QUESTION 91
You are creating a workflow in UCS Director's Workflow Designer.
Which of the following are not predefined workflow tasks? (Choose three.)

A. completed (failed)
B. completed
C. start
D. completed (success)
E. start (failed)
F. start (success)

Correct Answer: BEF

Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, the following Cisco Unified Computing System (UCS) Director's Workflow Designer tasks are not predefined: ▪ Completed
▪ Start (Failed)
▪ Start (success)
Workflows determine the order in which tasks that are designed to automate complex IT operations are performed. Workflow Designer allows administrators to create workflows that can then be automated by using UCS Director's orchestrator. By default, the following tasks are predefined when a workflow is created: ▪ Completed (failed)
▪ Completed (success)
▪ Start
The start task is the beginning of the workflow. The completed (failed) task represents the end of a workflow when the desired result could not be achieved. The completed (success) task represents a successfully completed workflow. Each task in a workflow processes input and produces output that is sent to the next task in the workflow. In addition, each task contains an On Success event and an On Failure event that can be used to determine which task should be performed next based on whether the task could be successfully completed.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 17: Understanding and Troubleshooting UCSD Workflows, Creating Workflows, pp. 641-645

QUESTION 92
Which of the following primary elements of a tenant defines the MAC address space?

A. a contract
B. a filter
C. a bridge domainD. an EPG

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

A bridge domain is the primary element of a tenant that defines the Media Access Control (MAC) address space. Bridge domains are logical Layer 2 forwarding configurations within a Cisco Application Centric Infrastructure (ACI) fabric that use switched virtual interfaces (SVIs) for gateways and can be configured to span multiple physical devices. In this respect, bridge domains are similar to virtual local area networks (VLANs). However, the purpose of a bridge domain is to define the MAC address space and flood domain.
Tenants are containers that can be used to represent organizations, domains, or specific groupings of information. Typically, tenants are configured to ensure that different policy types are isolated from each other, similar to user groups or roles in a role-based access control (RBAC) environment.
An endpoint group (EPG) is a primary element of a tenant; however, an EPG does not define the MAC address space. EPGs are logical groupings of endpoints that provide the same application or components of an application. For example, a collection of Hypertext Transfer Protocol Secure (HTTPS) servers could be logically grouped into an EPG labeled WEB. EPGs are typically collected within application profiles. EPGs can communicate with other EPGs by using contracts. A contract is a primary element of a tenant; however, a contract does not define the MAC address space. Contracts are policy objects that define how EPGs communicate with each other. There are three types of contracts that can be applied in an ACI fabric: ▪ Regular — applies filters to matching traffic and typically follows taboo contracts
▪ Taboo — denies and logs matching traffic
▪ Out-of-Band (OOB) - applies to OOB traffic from the management tenant
A filter is a primary element of a tenant; however, a filter does not define the MAC address space. Filters are low-level ACI objects that help define EPG contracts. Filters operate at Layer 2, Layer 3, and Layer 4 of the Open Systems Interconnection (OSI) networking model.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 13: ACI Logical Model and Policy Framework, ACI Logical Constructs, pp. 488-492
Cisco: Cisco APIC Basic Configuration Guide, Release 1.x: Tenants

QUESTION 93
Which of the following FCoE switch port types is typically connected to a VN port?

A. another VN port
B. a VF port
C. a SPAN port
D. a VE port

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

A Fibre Channel over Ethernet (FCoE) virtual fabric (VF) interface port type is typically connected to a virtual node (VN) port. FCoE is used in data centers to encapsulate Fibre Channel (FC) over an Ethernet network. This encapsulation enables the FC protocol to communicate over 10 gigabit-per-second (Gbps) Ethernet. There are two types of FCoE switch ports: a VF port and a virtual edge (VE) port.
An FCoE VF port typically connects to an end host. A VN port is a port on an end host, such as a host bus adapter (HBA) port. If the end host is connected to an Ethernet network that is configured with virtual local area networks (VLANs), the Spanning Tree Protocol (STP) configuration on the Ethernet fabric might require extra attention, especially if the Ethernet fabric is not using Per-VLAN Spanning Tree Plus (PVST+). A proper STP configuration on the Ethernet fabric prevents the Ethernet topology from affecting storage area network (SAN) traffic.
An FCoE VE port typically connects to a port on another FC forwarder (FCF). STP does not operate on VE ports, because these ports typically connect two FCFs. FC does not require switching loop prevention, because FCFs have no concept of switching loops. VE ports typically default to trunk mode.
A switched port analyzer (SPAN) port, which is also known as a mirroring port, is a type of port that is used to collect copies of packets transmitted over another port, over a given device, or over a network. In an FCoE configuration, a SPAN destination port can be either an FC interface or an Ethernet interface. SPAN source ports, on the other hand, can be FC interfaces, virtual FC (vFC) interfaces, a virtual SAN (vSAN), a VLAN, an Ethernet interface, a port channel interface, or a SAN port channel interface.

Reference:

Cisco: Fibre Channel over Ethernet Operations: FCoE and Spanning Tree Protocol Considerations

QUESTION 94
Which OTV failure isolation feature maps a MAC address from a remote data center to that remote data center's join interface IP address?

A. unknown unicast handling
B. broadcast policy control
C. ARP optimization
D. STP isolation

Correct Answer: A
Section: (none)
Explanation
Explanation/Reference:
Explanation:

Overlay Transport Virtualization (OTV)'s unknown unicast handling feature maps a Media Access Control (MAC) address from a remote data center to that data center's join interface Internet Protocol (IP) address. The join interface is the Open Systems Interconnection (OSI) networking model Layer 3 interface on which OTV is deployed and that is capable of discovering other OTV-enabled data centers. OTV's unicast traffic handling feature advertises MAC addresses between data centers. Because of this, advertisement and mapping process, OTV is capable of suppressing the transmission of unknown unicast traffic across data centers. OTV is a technology that extends Layer 2 networks across data centers. Similar to Ethernet over Multiprotocol Label Switching (EoMPLS), virtual private LAN services (VPLS), and dark fiber, OTV is intended to enable the extension of Layer 2 applications across large geographic distances. It is typically deployed on

data center edge devices. OTV uses MAC address routing to ensure that Layer 2 reachability information can be transmitted between data centers. Unlike other Layer 2 extension technologies, OTV's MAC address routing takes place at the control plane level. Other technologies rely on data plane forwarding and flooding. OTV's Spanning Tree Protocol (STP) isolation feature ensures that bridge protocol data units (BPDUs) are not forwarded across the overlay. This feature is enabled by default. Because BPDUs are suppressed, each data center's STP domain remains independent of other data centers connected to the overlay. This prevents data centers from accidentally creating Layer 2 loops over OTV.
Broadcast policy control and Address Resolution Protocol (ARP) optimization are OTV features that reduce traffic between OTV-connected data centers. Each of these features reduces the amount of broadcast traffic that is flooded between data centers.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 5: Data Center Overlay Networks, Failure Isolation, pp. 172-173

QUESTION 95
Which of the following formats are not used by a REST API to produce output? (Choose three.)

A. HTTP
B. JSON
C. CSV
D. XML
E. HTML

Correct Answer: ACE
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Representational state transfer (REST) does not produce output in comma separated values (CSV) format, Hypertext Markup Language (HTML) format, or Hypertext Transfer Protocol (HTTP) format. Instead, REST produces output in JavaScript Object Notation (JSON) or Extensible Markup Language (XML) format.
REST is an Application Programming Interface (API) architecture that uses HTTP or HTTP Secure (HTTPS) to enable external resources to access and make use of programmatic methods that are exposed by the API. For example, a web application that retrieves user product reviews from an online marketplace for display on third-party websites might obtain those reviews by using methods provided in an API that is developed and maintained by that marketplace. The JSON or XML output that is returned by the API is parsed by the third-party website for display.
HTTP is the Open Systems Interconnection (OSI) networking model Application layer protocol that is used to transfer information from a web server to a web browser. A REST API uses HTTP to transmit requests for information to a web server, which is not the same as producing the formatted output that is returned from the server.
Although HTML is similar to XML, which uses tags like HTML does, XML requires a strict syntax and is typically used to structure data, not format and render data in a web browser. HTML, on the other hand, is designed to inform a web browser about how given information should be displayed.

The CSV format is a common tabular format that is supported by spreadsheet applications and other business reporting applications. CSV files are plain-text files that segregate the fields of a table by using a combination of quotation marks, symbolic delimiters such as a comma or a semicolon, and line breaks.

Reference:

Cisco: Cisco APIC REST API Configuration Guide: About the REST API

QUESTION 96
Which of the following default user roles are common to both the Nexus 5000 switch and the Nexus 7000 switch? (Choose two.)

A. network-admin
B. san-admin
C. network-operator
D. vdc-admin
E. vdc-operator

Correct Answer: AC
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Both the network-admin user role and the network-operator user role are common default user roles to both the Nexus 5000 switch and the Nexus 7000 switch. However, the privileges associated with these roles are different between the switch models because the Nexus 5000 switch does not support the virtual device context (VDC) feature. A VDC is a virtual switch.
Cisco Nexus switches use role-based access control (RBAC) to assign management privileges to a given user. The Nexus 7000 switch is capable of supporting the VDC feature. By default, a Nexus 7000 switch is configured with the following user roles: ▪ network-admin — has read and write access to all VDCs on the switch ▪ network-operator — has read-only access to all the VDCs on the switch ▪ vdc-admin — has read and write access to a specific VDC on the switch ▪ vdc-operator — has read-only access to a specific VDC on the switch
Unlike the Nexus 7000 switch, a Nexus 5000 switch does not support the VDC feature. By default, a Nexus 5000 switch is configured with the following user roles: ▪ network-admin — has complete read and write access to the switch ▪ network-operator — has read-only access to the switch ▪ san-admin — has complete read and write access to Fibre Channel (FC) and FC over Ethernet (FCoE) by using Simple Network Management Protocol (SNMP) or the command-line interface (CLI)

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 2: Management and Monitoring of Cisco Nexus Devices, User Roles, p. 71

QUESTION 97
Which of the following default Cisco UCS Manager user roles is automatically assigned to the default admin account?

A. Administrator
B. Operations
C. AAA Administrator
D. Read-Only
E. Network Administrator

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

The Cisco Unified Computing System (UCS) default user role of Administrator is automatically assigned to the default admin account. The Cisco UCS System is installed with several default user roles for its role-based access control (RBAC) system. User accounts in Cisco UCS Manager are assigned one of the default roles or a custom role that is defined by an administrator. These roles define the access privileges for each account. Cisco Manager is web-based software that can be used to manage a single UCS domain. Administrators have read-write access to the entire system. Although the Administrator role can be assigned to other user accounts, this role cannot be removed from the default admin account.
The AAA Administrator user role is not automatically assigned to the default admin account. The AAA Administrator user role has read-write access to users, roles, and the system's authentication, authorization, and accounting (AAA) configuration. AAA Administrators can read the rest of the system but cannot write to it. The Network Administrator user role is not automatically assigned to the default admin account. The Network Administrator user role has read-write access to the fabric interconnect and network security. However, this role has only read access to the rest of the system.
The Operations user role is not automatically assigned to the default admin account. The Operations user role has read-write access to system logs and read access to the rest of the system. An Operations user can read and write to Syslog servers and faults.
The Read-Only user role is not automatically assigned to the default admin account. The Read-Only user role has read access to the configuration. In addition, the Read-Only user role cannot modify the system state. There are five other default Cisco UCS Manager user roles: ▪ Facility Manager — has read-write access to power management
▪ Server Equipment Administrator — has read-write access to physical server operations
▪ Server Profile Administrator — has read-write access to logical server operations
▪ Server Security Administrator — has read-write access to security operations ▪
Storage Administrator — has read-write access to storage operations

Reference:

Cisco: Configuring Role-Based Access Control: Default User Roles

QUESTION 98
You want to configure user authentication methods and user roles in the Cisco UCS Manager GUI.
Which of the following Navigation pane tables should you click?

A. **LAN**
B. **Admin**
C. **Equipment**
D. **Servers**
E. **SAN**
F. **VM**

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

You should click the **Admin** tab in the Cisco Unified Computing System (UCS) Manager graphical user interface (GUI) if you want to configure authentication methods and user roles. Cisco UCS Manager GUI is a Java application. The main area of the GUI is divided into a Navigation pane and a work area. The selections you click in the Navigation pane determine the information and configuration fields that are displayed in the work area.
The Navigation pane of the Cisco UCS Manager GUI contains six tabs in a row across the top of the pane. To configure or view a given element of Cisco UCS Manager, you should first click the Navigation pane tab appropriate for that element. The Navigation pane contains all of the following tabs: ▪ The **Equipment** tab
▪ The **Servers** tab
▪ The **LAN** tab
▪ The **SAN** tab
▪ The **VM** tab
▪ The **Admin** tab
The **Equipment** tab can be used to display an inventory of the UCS domain. This information includes color-coded fault indicators, such as a red, yellow, or orange rectangle. If a device has a fault, one of these indicators will appear around the name of the device on the **Equipment** tab. The **Equipment** tab contains four nodes: **Equipment**, **Chassis**, **Rack-mounts**, and **Fabric Interconnects**. Selected nodes contain information specific to the devices indicated by the node name. For example, the **Fabric Interconnects** node contains information about expansion modules, fans, and power supply units (PSUs) connected to the domain's fabric interconnects.
The **Servers** tab can be used to modify server-specific configurations, such as policies, profiles, and universally unique identifier (UUID) pools. The **Servers** tab contains six nodes: **Servers**, **Service Profiles**, **Service Profile Templates**, **Policies**, **Pools**, and **Schedules**. Selected nodes contain information specific to the server configurations indicated by the node name. For example, the **Policies** node allows the configuration of policies related to server adapters, server firmware, and other components.

The **LAN** tab can be used to configure local area network (LAN) components, such as Quality of Service (QoS) classes, virtual LANs (VLANs), and flow control policies. The **LAN** tab contains seven nodes: **LAN Cloud**, **Appliances**, **Internal LAN**, **Policies**, **Pools**, **Traffic Monitoring Sessions**, and **Netflow Monitoring**. Selected nodes contain information specific to the LAN component indicated by the node name. For example, the **Pools** node allows the configuration of both IP address pools and MAC address pools that have been defined for a LAN.
The **SAN** tab can be used to configure storage area network (SAN) components, such as virtual SANs (vSANs), and World Wide Name (WWN) pools. The **SAN** tab contains six nodes: **SAN**, **SAN Cloud**, **Storage Cloud, Policies**, **Pools**, and **Traffic Monitoring Sessions**. Selected nodes contain information specific to the SAN component indicated by the node name. For example, the **SAN** node allows the configuration of SAN uplinks, Fibre Channel (FC) address assignments, and vSANs.
The **VM** tab can be used to configure virtual machine-fabric extender (VM-FEX) for UCS domain servers that are equipped with virtual interface cards (VICs). The **VM** tab contains seven nodes: **All**, Clusters, **Fabric Network Sets**, **Port Profiles**, **VM Networks**, **Microsoft**, and **VMware**. Selected nodes contain information specific to the VM component indicated by the node name. For example, the **VMware** node can be used to configure Cisco UCS Manager connections to VMware vCenter.
The **Admin** tab can be used to configure system-wide settings that must be configured by an administrator or viewed by a security administrator. The **Admin** tab contains 10 nodes:

- **All**
- **Faults, Events and Audit Log**
- **User Management**
- **Key Management**
- **Communication Management**
- **Stats Management**
- **Time Zone Management**
- **Capability Catalog**
- **Management Extension**
- **License Management**

Selected nodes contain information specific to the administrative component indicated by the node name. For example, the **User Management** node allows the configuration of authentication methods and user roles as well as remote access methods.

Reference:

Cisco: Overview of Cisco UCS Manager GUI: Navigation Pane

QUESTION 99
Which of the following best describes a policy repository?

A. a controller for distributing, tracking, and updating policies
B. collection of endpoints with identical requirements
C. a collection of rules applied to endpoints, existing or hypothetical
D. a collection of endpoints already known to the Cisco ACI

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, a policy repository can be best described as a collection of rules applied to endpoints, existing or hypothetical. A policy repository is the first of two logically coupled components of an endpoint policy; it is typically centralized in the Cisco Application Policy Infrastructure Controller (APIC) or in leaf nodes in the Application Centric Infrastructure (ACI) fabric. The policy repository component of an endpoint policy stores rules that can be applied to existing endpoints, planned endpoints, or deleted endpoints. Endpoint policies are used to define how endpoint groups (EPGs) communicate with each other.
An endpoint registry is a collection of endpoints already known to the Cisco ACI. The endpoint registry is the second of two logically coupled components of an endpoint policy; it stores the operational state of each endpoint. In Cisco ACI, the endpoint registry is stored in a distributed database within the ACI fabric. An EPG is a collection of endpoints with identical requirements; it is a primary element of a tenant. EPGs are logical groupings of endpoints that provide the same application or components of an application. For example, a collection of Hypertext Transfer Protocol Secure (HTTPS) servers could be logically grouped into an EPG labeled WEB. EPGs are typically collected within application profiles. EPGs can communicate with other EPGs by using contracts.
The Cisco APIC is a controller for distributing, tracking, and updating policies. In other words, it is a central component of a Cisco ACI fabric that is used for both automation and management.

Reference:

Cisco: The Cisco Application Policy Infrastructure Controller: Endpoints and Policy Control

QUESTION 100
Which of the following statements best describes vPC peer keepalive links?

A. There can be only two peers per domain.
B. They synchronize the control plane and the data plane.
C. They monitor the status of vPC peers.
D. They synchronize the state between two vPC peers.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A virtual port channel (vPC) peer keepalive link monitors the status of vPC peers. The peer keepalive link operates at Layer 3 of the Open Systems Interconnection (OSI) networking model; it is used to ensure that vPC switches are capable of determining whether a vPC domain peer has failed. Peer keepalive links can be

configured to operate in any virtual routing and forwarding (VRF) instance, including the management VRF. Each vPC peer keepalive link is configured with the remote peer's IP address as its destination IP address and the local peer's IP address as its source address. Peer keepalive links must be trunk links.
There can be only two peers, or switches, per vPC domain. A vPC enables you to bundle ports from two peers, which form a domain, into a single OSI Layer 2 port channel. Similar to a normal port channel, a vPC bundles multiple switch ports into a single high-speed trunk port. A single vPC domain cannot contain ports from more than two switches. For ports on two switches to successfully form a vPC domain, all the following must be true: ▪ The vPC feature must be enabled on both switches.
▪ The vPC domain ID must be the same on both switches.
▪ The peer keepalive link must be configured and must be 10 gigabits per second (Gbps) or more. ▪
The vPC number must be the same on both switches.
A vPC peer link, not a vPC peer keepalive link, synchronizes the state between two vPC peers. A vPC peer link is typically comprised of a port channel made up of two physical ports on each switch. This link synchronizes Media Access Control (MAC) address tables between switches and serves as a transport for data plane traffic. Bridge protocol data unit (BPDU) and Link Aggregation Control Protocol (LACP) packets are also forwarded to the second peer over this link, which causes the vPC peers to appear to be a single control plane.
Cisco Fabric Services, not a vPC peer keepalive link, synchronizes the control plane and the data plane. Cisco Fabric Services is a messaging protocol that operates between vPC peers. Control plane and data plane information is synchronized over the vPC peer link.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 1: Data Center Networking, Components of vPC, pp. 22-24

QUESTION 101
Which of the following FabricPath components provides access layer connectivity?

A. a leaf switch
B. the APIC
C. a spine switch
D. a CE network

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, a leaf switch is the Cisco FabricPath component that provides access layer connectivity. Cisco FabricPath is a means of constructing a scalable Open Systems Interconnection (OSI) networking model Layer 2 network from both Layer 2 and Layer 3 components. End hosts and classic Ethernet (CE) networks are typically directly connected to leaf switches by using edge ports.

Spine switches do not provide access layer connectivity. Spine switches are the Cisco FabricPath component that form the backbone of the FabricPath's switching fabric. Typically, leaf switches are connected to every spine switch along the backbone so that the spine switches provide connectivity between the leaf switches. Leaf switches connect to spine switches by using core ports.
The Cisco Application Policy Infrastructure Controller (APIC) is a means of managing the Cisco Application Centric Infrastructure (ACI). A Cisco ACI architecture requires both the APIC and the spine switches and leaf switches of FabricPath to complete the architecture. The APIC communicates with the spine and leaf nodes and provides policy distribution as well as centralized management.
A CE network, which is a traditional Ethernet network that uses Spanning Tree Protocol (STP) and transparent bridging, is not technically part of Cisco FabricPath. However, Cisco Nexus switches can connect to a CE network and Cisco FabricPath simultaneously.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 1: Data Center Networking, Components of FabricPath, pp. 32-34

QUESTION 102
Which of the following is Cisco software that can be used to manage multiple UCS domains across geographical boundaries?

A. Cisco UCS Director
B. Cisco IMC Supervisor
C. Cisco UCS Manager
D. Cisco UCS Central

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Cisco Unified Computing System (UCS) Central is software that can be used to manage multiple UCS domains, including domains that are separated by geographical boundaries. Cisco UCS Central can be used to deploy standardized configurations and policies from a central virtual machine (VM).
Cisco UCS Director is software that can automate actions and be used to construct a private cloud. Cisco UCS Director creates a basic Infrastructure as a Service (IaaS) framework by using hardware abstraction to convert hardware and software into programmable actions that can then be combined into an automated custom workflow. Thus Cisco UCS Director enables administrators to construct a private cloud in which they can automate and orchestrate both physical and virtual components of a data center. Cisco UCS Director is typically accessed by using a web-based interface.
Cisco UCS Manager is web-based software that can be used to manage a single UCS domain. The software is typically embedded in Cisco UCS fabric interconnects rather than installed in a VM or on separate physical servers.
Cisco Integrated Management Controller (IMC) Supervisor is software that can be used to centrally manage multiple standalone Cisco C-Series and E-Series servers. The servers need not be located at the same site. Cisco IMC Supervisor uses a web-based interface and is typically deployed as a downloadable virtual application.

Reference:
Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 7: Cisco UCS Architecture, Cisco UCS Software, pp. 274-277

QUESTION 103
Which of the following is true of the default VRF on a Cisco router?

A. It includes only the mgmt 0 interface.
B. No routing protocols are allowed to run there.
C. It is similar to a router's global routing table.
D. It is used only for management.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

A Cisco router's default virtual routing and forwarding (VRF) instance is similar to a router's global routing table. The default VRF includes all Layer 3 interfaces until you assign those interfaces to another VRF. Similarly, the default VRF runs any routing protocols that are configured unless those routing protocols are assigned to another VRF. All **show** and **exec** commands that are issued in the default VRF apply to the default routing context.
VRFs are used to logically separate Open Systems Interconnection (OSI) networking model Layer 3 networks. Therefore, it is possible to have overlapping Internet Protocol version 4 (IPv4) or Internet Protocol version 6 (IPv6) addresses in environments that contain multiple tenants. However, an interface that has been assigned to a given VRF cannot be simultaneously assigned to another VRF. The address space, routing process, and forwarding table that are used within a VRF are local to that VRF. By default, a Cisco router is configured with two VRFs: the management VRF and the default VRF.
The management VRF, not the default VRF, is used only for management. No routing protocols are allowed to run in the management VRF. All routing is static. The management VRF includes only the mgmt 0 interface, which cannot be assigned to any other VRF. However, the mgmt 0 interface is shared among virtual device contexts (VDCs).

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 6: Virtualizing Cisco Network Devices, Describing Layer 3 Virtualization Within VDCs, pp. 204-206

QUESTION 104
You are configuring a service profile for a Cisco UCS server that contains two physical HBAs and no NICs.
How many vNICs can be configured for this server?

A. four

B. eight
C. six
D. none
E. two

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

No virtual network interface cards (vNICs) can be configured for a Cisco Unified Computing System (UCS) server that contains two physical host bus adapters (HBAs) and no network interface cards (NICs). Cisco UCS server can be configured with the number of vNICs that corresponds to available physical NICs on the server. Similarly, a Cisco UCS can be configured with the number of virtual HBAs (vHBAs) that corresponds to available physical HBAs on the physical adapters that are installed in the device. In this scenario, it is possible to configure the server with two VHBAs.
It is not possible to configure two vNICs for the Cisco UCS server in this scenario. In order to configure two NICs, the server would need to be configured with one or more adapters that each contains one or more physical NICs. For example, a Cisco converged network adapter typically contains two physical ports. Therefore, a Cisco UCS server that is configured with a single converged network adapter could be configured with two NICs. A converged network adapter is a single unit that combines a physical HBA and a physical Ethernet NIC. A Cisco converged network adapter typically contains two of these types of ports.
It is not possible to configure four vNICs for the Cisco UCS server in this scenario. To configure four vNICs, the UCS server would need to be configured with adapters that contain up to four physical NICs.
It is not possible to configure eight VNICs for the Cisco UCS server in this scenario. To configure eight vNICs, the UCS server would need to be configured with adapters that contain up to eight physical NICs.

Reference:

Cisco: Overview of Cisco Unified Computing System: Configuration through Service Profiles.

QUESTION 105
Which of the following Cisco UCS servers is most likely to require a special version of Cisco UCS Manager?

A. C-Series servers
B. B-Series servers
C. E-Series servers
D. UCS Mini servers

Correct Answer: D
Section: (none)

Explanation
Explanation/Reference:
Explanation:

Of the available options, Cisco Unified Computing System (UCS) Mini servers are most likely to require a special version of Cisco UCS Manager. Cisco UCS Mini servers are a compact integration of a Cisco UCS 5108 blade chassis, Cisco UCS 6324 Fabric Interconnects, and Cisco UCS Manager. Unlike other Cisco UCS servers, the UCS Mini server requires a special version of Cisco UCS Manager for management. It is possible to connect Cisco UCS C-Series servers to Cisco UCS Mini servers in order to expand their abilities.
Cisco UCS C-Series servers do not require a UCS fabric, nor do they require Cisco UCS Manager. Cisco UCS C-Series servers are rack-mount standalone servers that can operate either with or without integration with Cisco UCS Manager. For administrators who are more familiar with traditional Ethernet networks than UCS Fabric Interconnect, C-Series servers will most likely be simpler to deploy and feel more familiar than other Cisco UCS server products.
Cisco UCS B-Series servers require a UCS fabric and use a standard version of Cisco UCS Manager. Cisco UCS B-Series servers are blade servers that are installed in a UCS blade chassis. These blade servers can connect only to Cisco UCS Fabric Interconnect, not directly to a traditional Ethernet network. Blade servers in a chassis are typically hot-swappable, unlike the components of a rack-mount server. Therefore, blade server configurations are less likely to result in prolonged downtime if hardware fails.
Cisco UCS E-Series servers are blade servers that can use a standard version of Cisco UCS Manager. However, Cisco UCS E-Series servers have similar capabilities to the standalone C-Series servers and do not require connectivity to a UCS fabric. In a small office environment, Cisco UCS E-Series servers are capable of providing the network connectivity and capabilities of a C-Series server along with the availability of B-Series servers.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 7: Cisco UCS Architecture, Cisco UCS Mini, p. 272

QUESTION 106
Which of the following are true about Cisco UCS Director workflow templates? (Choose three.)

A. They contain task names, a workflow structure, and input names.
B. They can be executed like normal workflows.
C. They can be used to schedule workflows.
D. They can be used to instantiate new workflows.
E. Some templates are predefined in the system.

Correct Answer: ADE
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Cisco Unified Computing System (UCS) Director workflow templates can be used to instantiate new workflows. Workflow templates are used as blueprints to create new workflows. In addition, you can create a new workflow and then export that workflow as a template in order to create other workflows built around the same tasks.
Cisco UCS Director workflow templates contain task names, a workflow structure, and input names. Task names are literally the string names of the tasks that are to be executed in the workflow. The tasks themselves are already built into the system in which the workflow template is defined. A workflow structure defines how the tasks are connected in the series. Input names define the details about the variables that are used to store the workflow's user input.
Some Cisco UCS Director workflow templates are predefined in the system. Predefined workflow templates can be accessed by using the **Workflow Templates** tab in Cisco UCS Director Orchestration.
Cisco UCS Director workflow templates cannot be executed like normal workflows. In addition, workflow templates cannot be used to schedule workflows. Cisco UCS Director is a resource automation and orchestration tool. UCS Director workflows are task sequences that accept user input and then automatically perform a series of tasks to complete a complex operation. Workflow templates simplify the creation of workflows based on a set of predefined tasks and therefore cannot themselves be executed. Workflows can be executed in several ways, including by using service requests, the Execute Now action, a virtual machine (VM) Action policy, a trigger, a schedule, or the rollback feature.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 17: Understanding and Troubleshooting UCSD Workflows, Creating Workflows, pp. 641-645

QUESTION 107
Which of the following is least likely to be used to construct or access a cloud-based API?

A. a SOAP API
B. GraphQL
C. the Java API
D. a REST API

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, the Java Application Programming Interface (API) is least likely to be used to construct or access a cloud-based API. The Java API is typically accessed by Java applications that are running in the Java virtual machine (VM), which is the Java component that executes compiled Java programs. The Java API is a collection of Java classes that developers can use for data collection or to build interfaces. Open APIs can be used to enable services such as billing automation and centralized management of cloud infrastructure.

Representational state transfer (REST) is an API architecture that uses Hypertext Transfer Protocol (HTTP) or HTTP Secure (HTTPS) to enable external resources to access and make use of programmatic methods that are exposed by the API. Therefore, it is possible to construct or access a cloud-based API from REST. For example, a web application that retrieves user product reviews from an online marketplace for display on third-party websites might obtain those reviews by using methods provided in an API that is developed and maintained by that marketplace. REST APIs can return data in Extensible Markup Language (XML) format or in JavaScript Object Notation (JSON) format.
Simple Object Access Protocol (SOAP) APIs can be used to construct or access cloud-based APIs. SOAP is an older API messaging protocol that uses HTTP and XML to enable communication between devices. SOAP APIs are typically more resource-intensive than REST APIs and, therefore, slower. Unlike REST APIs, SOAP APIs do not return JSON-formatted output.
Graph Query Language (GraphQL) can be used to access cloud-based APIs. GraphQL is an API query language and a runtime that is intended to lower the burden of making multiple API calls to obtain a single set of data. For example, data that requires three or four HTTP GET requests when constructed from a standard REST API might take only one request when using GraphQL. Similar to REST API, GraphQL output is in JSON format. Although GraphQL can use HTTP or HTTPS, it is not limited to those protocols.

Reference:

Cisco Press CCNA Data Center DCICT 200-155 Official Cert Guide, Chapter 15: Cloud Computing, Application Programming Interfaces, pp. 577-579

QUESTION 108
You connect an FC link to unified port 1/31 on a new Cisco Nexus 5548UP switch and receive the following error messages:

```
ERROR: Ethernet range starts from first port of the module
ERROR: FC range should end on last port of the module
```

Which of the following is most likely true?

A. The link you connected should have been to unified port 1/32.
B. The link you connected was actually Ethernet, not FC.
C. The link you connected should have been to unified port 1/17.
D. The link you connected was FC but should have been Ethernet.
E. The link you connected should have been to unified port 1/1.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Of the available choices, it is most likely that the link you connected should have been to unified port 1/32 on the Cisco Nexus 5548UP switch in this scenario. Unified ports are ports that can be either Ethernet or Fibre Channel (FC) ports. Cisco Nexus switches that support unified ports require that Ethernet links be connected from the beginning of the port range forward and that FC links be connected from the end of the port range backward. Slot 1 of a Cisco Nexus 5548UP

switch contains 32 unified ports. Therefore, in this scenario, the FC link on the new switch would need to be connected to port 1/32 before any other FC links could be connected in the FC range.
As a further example, if you were to connect a single Ethernet link and a single FC link to slot 1 of a Cisco Nexus 5548UP switch that has no other links connected, you must connect the Ethernet link to port 1/1 and the FC link to port 1/32. If you were to then connect an additional Ethernet link to slot 1, that link must be connected to port 1/2. If you were to connect an additional FC link to slot 1, that link must be connected to port 1/31. Link additions should continue in that way until the Ethernet range ends where the FC range begins and no more ports are available in the slot.
The link in this scenario should not have been connected to unified port 1/17. Unified port 1/17 on a Cisco Nexus 5548UP switch falls in the middle of the 32-port range of slot 1. This port could be the end of an Ethernet port range if it were preceded by 16 other Ethernet links or the beginning of an FC port range if it were succeeded by 15 other FC links.
The link in this scenario should not have been connected to unified port 1/1. In addition, there is nothing in this scenario to indicate that the link was or should have been an Ethernet link.

Reference:

Cisco: Configuring Layer 2 Interfaces: Information About Unified Ports